arXiv:1611.09014v2 [cs.AI] 29 Nov 2016

1

Blocking and Other Enhancements for Bottom-Up Model Generation Methods

Peter Baumgartner Renate A. Schmidt

23 November 2016

Abstract

Model generation is a problem complementary to theorem proving and is important for fault analysis and debugging of formal specifications of, for example, security protocols, programs and terminological definitions. This paper discusses several ways of enhancing the paradigm of bottom-up model generation. The two main contributions are new, generalized blocking techniques and a new range-restriction transformation. The blocking techniques are based on simple transformations of the input set together with standard equality reasoning and redundancy elimination techniques. These provide general methods for finding small, finite models. The range-restriction transformation refines existing transformations to range-restricted clauses by carefully limiting the creation of domain terms. All possible combinations of the introduced techniques and a classical range-restriction technique were tested on the clausal problems of the TPTP Version 6.0.0 with an implementation based on the SPASS theorem prover using a hyperresolution-like refinement. Unrestricted domain blocking gave best results for satisfiable problems showing it is a powerful technique indispensable for bottom-up model generation methods. Both in combination with the new range-restricting transformation, and the classical range-restricting transformation, good results have been obtained. Limiting the creation of terms during the inference process by using the new range restricting transformation has paid off, especially when using it together with a shifting transformation. The experimental results also show that classical range restriction with unrestricted blocking provides a useful complementary method. Overall, the results showed bottom-up model generation methods were good for disproving theorems and generating models for satisfiable problems, but less efficient than SPASS in auto mode for unsatisfiable problems.

1 Introduction

The bottom-up model generation (BUMG) paradigm encompasses a wide family of calculi and proof procedures that explicitly try to construct a model of a given clause set by reading clauses as rules and applying them in a bottom-up way until completion. For instance, variants of hyperresolution and grounding

*P. Baumgartner, Data61, CSIRO, Australia
R. A. Schmidt, School of Computer Science, The University of Manchester, UK

tableau calculi belong to this family. BUMG methods have been known for a long time to useful for proving theorems, comparably little effort has however been undertaken to exploit them for the dual task, namely, computing models for satisfiable problems. This is somewhat surprising, as computing models is recognized as being important in software engineering, model checking, and other applications for fault analysis and debugging of logical specifications.

One of the contributions of the paper is the introduction to first-order logic of blocking techniques partially inspired by techniques already successfully used in description and modal logic tableau-based theorem proving [Hustadt and Schmidt, 1999, Schmidt and Tishkovsky, 2007, Baader and Sattler, 2001]. We adapt and generalize these blocking techniques to full first-order logic. Blocking is an important technique for turning tableau systems into decision procedures for modal and description logics. Though different blocking techniques exist, and not all modal and description logic tableau systems are designed to return models, blocking is essentially a mechanism for systematically merging terms in order to find finite models.

In our approach blocking is encoded on the clausal level and is combined with standard resolution techniques, the idea being that with a suitable prover small, finite models are constructed and can be easily read off from the derived clauses. Our blocking techniques are generic and pose no restrictions on the logic they can be used for. They can even be used for undecidable logics. We introduce four different blocking techniques. The main idea of our blocking techniques is that clauses are added to the input problem which lead in the derivation to splittable clauses causing terms in the partially constructed models to be merged. The difference between the four techniques is how restrictive blocking is. With *unrestricted domain blocking* domain minimal models can be generated. With *subterm domain blocking* or *subterm predicate blocking* larger models are produced because two terms are only merged if one is a subterm of the other. With *unrestricted predicate blocking* and *subterm predicate blocking* two terms are merged if they both belong to the extension of a unary predicate symbol, the intention being that less constrained, finite model can be found.

The second contribution of the paper is a refinement of the well-known 'transformation to range-restricted form' as introduced in the eighties by Manthey and Bry [1988] in the context of the SATCHMO prover and later improved, for example, by Baumgartner et al. [1997]. These range-restricting transformations have the disadvantage that they generally force BUMG methods to enumerate the entire Herbrand universe and are therefore non-terminating except in the simplest cases. One solution is to combine classical range-restriction transformations with blocking techniques. Another solution, presented in this paper, is to modify the range-restricting transformation so that new terms are created only when needed. Our method extends and combines the range-restricting transformation introduced in Schmidt and Hustadt [2005] for reducing first-order formulae and clauses into range-restricted clauses, which was used to develop general-purpose resolution decision procedures for the Bernays-Schönfinkel class.

Other methods for model computation can be classified as methods that directly search for a finite model, such as the extended PUHR tableau method of Bry and Torge [1998], the method of Bezem [2005] and the methods in the SEM-family [Slaney, 1992, Zhang, 1995, McCune, 2003]. In contrast, MACE-style model builders such as, for example, the methods of Claessen and Sörensson [2003] and McCune [1994] reduce model search to testing of propositional satisfi-

ability. Being based on a translation, the MACE-style approach is conceptually related, but different to our approach. Both SEM- and MACE-style methods search for finite models, by essentially searching the space of interpretations with domain sizes $1, 2, \ldots$, in increasing order, until a model is found.

Our method operates significantly differently, as it is *not* parameterized by a domain size. Consequently, there is no requirement for iterative deepening over the domain size, and the search for finite models works differently. This way, we can address a problem often found with models computed by these methods: from a pragmatic perspective, they tend to identify too many terms. For instance, for the two unit clauses $\mathsf{P}(\mathsf{a})$ and $\mathsf{Q}(\mathsf{b})$ there is a model that identifies a and b with the same object. Such models can be counter-intuitive, for instance, in a description logic setting, where unique names are often assumed, but not necessarily explicitly specified. Furthermore, logic programs are typically understood with respect to Herbrand semantics, and it is desirable to develop compatible model building techniques. We present transformations that are more careful at identifying objects than the methods mentioned and thus work closer to a Herbrand semantics.

The structure of the paper is as follows. Definitions of basic terminology and notation can be found in Section 2. In Section 3 we recall the characteristic properties of BUMG methods. The main part of the paper are Sections 4 to 9. Sections 4, 5 and 6 define new techniques for generating small models and generating them more efficiently. The techniques are based on a series of transformations including a refined range-restricting transformation (Section 4), instances of standard renaming and flattening (Section 5), and the introduction of blocking in various forms through amendments of the clause set and standard saturation-based equality reasoning (Section 6). Soundness and completeness of the blocking transformations and the combined transformations is shown in Section 7. One consequence of the results is a general decidability result of the Bernays-Schönfinkel class for all BUMG methods and related approaches. This is presented in Section 8. In Section 9 we present and discuss results of experiments carried out with our methods on clausal problems in the TPTP library.

This paper is an extended and improved version of Baumgartner and Schmidt [2006].

2 Basic Definitions

We use standard terminology from automated reasoning. We assume as given a signature $\Sigma = \Sigma_f \cup \Sigma_P$ of function symbols Σ_f (including constants) and predicate symbols Σ_P. As we are working (also) with equality, we assume Σ_P contains a distinguished binary predicate symbol $\approx$, which is used in infix form. Terms, atoms, literals and formulas over Σ and a given (denumerable) set of variables V are defined as usual.

A clause is a (finite) implicitly universally quantified disjunction of literals. We write clauses in a logic-programming style, that is, we write $H_1 \vee \cdots \vee H_m \leftarrow B_1 \wedge \cdots \wedge B_k$ rather than $H_1 \vee \cdots \vee H_m \vee \neg B_1 \vee \cdots \vee \neg B_k$, where $m, k \geq 0$. Each H_i is called a *head atom*, and each B_j is called a *body atom*. When writing expressions such as $H \vee \mathcal{H} \leftarrow B \wedge \mathcal{B}$ we mean any clause whose head literals are H and those in the disjunction of literals $\mathcal{H}$, and whose body literals are B and

those in the conjunction of literals $\mathcal{B}$. A *clause set* is a finite set of clauses.

A clause $\mathcal{H} \leftarrow \mathcal{B}$ is said to be *range-restricted* iff the body $\mathcal{B}$ contains all the variables in it. This means that a positive clause $\mathcal{H} \leftarrow \top$ is range-restricted only if it is a ground clause. A clause set is range-restricted iff it contains only range-restricted clauses.

For a given atom $P(t_1, \ldots, t_n)$ the terms $t_1, \ldots, t_n$ are also called the *top-level terms* of $P(t_1, \ldots, t_n)$ (P being $\approx$ is permitted). This notion generalizes to clause bodies, clause heads and clauses as expected. For example, for a clause $\mathcal{H} \leftarrow \mathcal{B}$ the top-level terms of its body $\mathcal{B}$ are exactly the top-level terms of its body atoms.

A *proper functional term* is a term that is neither a variable nor a constant.

A *(Herbrand) interpretation* I is a set of ground atoms, namely, those that are true in the interpretation. Satisfiability/validity in a Herbrand interpretation of ground literals, clauses, and clause sets is defined as usual. Also, as usual, a clause set stands semantically for the set of all its ground instances. We write $I \models F$ to denote that I satisfies F, where F is a ground literal or a (possibly non-ground) clause (set).

An *E-interpretation* is an interpretation that is also a congruence relation on the terms in the signature. If I is an interpretation, we denote by I^E the smallest congruence relation on the terms that includes I, which is an E-interpretation. An E-interpretation does not necessarily need to be a Herbrand-E-interpretation and is a standard first-order interpretation I such that $(I, \mu) \models s \approx t$ if and only if $(I, \mu)(s) = (I, \mu)(t)$ (where μ is a valuation, that is, a mapping from the variables to the domain $|I|$ of I). We say that I *E-satisfies* F iff $I^{\mathrm{E}} \models F$. Instead of $I^{\mathrm{E}} \models F$ we write $I \models_{\mathrm{E}} F$.

It is well-known that E-interpretations can be characterized by fixing the domain as the Herbrand universe and requiring that for every ground term t, $t \approx t \in I$, and for every ground atom A (including ground equations) the following is true: whenever $I \models A[s]$ and $I \models s \approx t$, then $I \models A[t]$.

Another characterization is to add to a given clause set M its equality axioms $\mathrm{EAX}(\Sigma_P \cup \Sigma_f)$, that is, the axioms expressing that $\approx$ is a congruence relation on the terms and atoms induced by the predicate symbols Σ_P and function symbols Σ_f occurring in M. It is well-known that M is E-satisfiable iff $M \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f)$ is satisfiable.

We work mostly, but not always, with Herbrand interpretations. If not, we always make this clear, and the interpretations considered then are first-order logic interpretations with domains that are (proper) subsets of the Herbrand universe of the clause set under consideration. Such interpretations are called *quasi-Herbrand interpretations*. When constructing such interpretations the requirement that function symbols are interpreted as total functions over their domain is not always trivially satisfied. For instance, in the presence of a constant a, a unary function symbol f, and the domain $\{a, f(a)\}$, say, one has to assign a value in the interpretation to every term. However $f(f(a))$, for instance, cannot be assigned to itself, as $f(f(a))$ is not contained in the domain.

3 BUMG Methods

Proof procedures based on model generation approaches establish the satisfiability of a problem by trying to build a model for the problem. In this paper

we are interested in bottom-up model generation approaches (BUMG). BUMG approaches use a forward reasoning approach where implications or clauses, $\mathcal{H} \leftarrow \mathcal{B}$, are read as rules and are repeatedly used to derive (instances of) $\mathcal{H}$ from (instances of) $\mathcal{B}$ until a completion is found.

The family of BUMG approaches includes many familiar calculi and proof procedures such as Smullyan type semantic tableaux [Smullyan, 1971], SATCHMO [Manthey and Bry, 1988, Geisler et al., 1997], positive unit hyperresolution (PUHR) tableaux [Bry and Yahya, 2000, Bry and Torge, 1998], the model generation theorem prover MGTP [Fujita et al., 1995] and hypertableaux [Baumgartner et al., 1996]. A well-established and widely known method for BUMG is hyperresolution [Robinson, 1965].

Hyperresolution consists of two inference rules, hyperresolution and factoring. The *hyperresolution rule* applies to a non-positive clause $H \leftarrow B_1 \wedge \ldots \wedge B_n$ ($n \neq 0$) and n positive clauses $C_1 \vee B'_1 \leftarrow \top$, ..., $C_n \vee B'_n \leftarrow \top$, and derives $(C_1 \vee \ldots \vee C_n \vee H)\sigma \leftarrow \top$, where σ is the most general unifier such that $B'_i\sigma = B_i\sigma$ for every $i \in \{1, \ldots, n\}$. The *factoring rule* derives the clause $(C \vee B)\sigma \leftarrow \top$ from a positive clause $C \vee B \vee B' \leftarrow \top$, where σ is the most general unifier of B and B'. On range-restricted clauses, when using hyperresolution, factoring amounts to the elimination of duplicate literals in positive clauses and is therefore optional when clauses are viewed as sets.

A crucial requirement for the effective use of blocking (considered later in Section 6) is support of equality reasoning (for example, ordered paramodulation, ordered rewriting or superposition [Bachmair and Ganzinger, 1998, Nieuwenhuis and Rubio, 2001]), in combination with simplification techniques based on orderings. We refer to Bachmair and Ganzinger [1998, 2001] for general notions of redundancy in saturation-based theorem proving approaches.

Our experiments show the splitting rule is useful for BUMG. For our blocking transformations, splitting on the positive part of (ground) clauses is in fact mandatory to make it effective. This type of splitting replaces the branch of a derivation containing a positive clause $C \vee D \leftarrow \top$, say, by two copies of the branch in which the clause is replaced by $C \leftarrow \top$ and $D \leftarrow \top$, respectively, provided that C and D do not share any variables. Most BUMG procedures support this splitting technique, in particular, the provers we have used do.

4 Range-Restricting Transformations

Existing transformations to range-restricted form follow Manthey and Bry [1988] (or are variations of it). The transformation can be defined by a procedure carrying out the following steps on a given set M of clauses.

(0) Initialization. Initially, let $\mathrm{crr}(M) := M$.

(1) Add a constant. Let dom be a 'fresh' unary predicate symbol not in Σ_P, and let c be some constant. Extend $\mathrm{crr}(M)$ by the clause

$$\mathsf{dom}(c) \leftarrow \top.$$

The constant c can be 'fresh' or belong to Σ_f.

(2) Range-restriction. For each clause $\mathcal{H} \leftarrow \mathcal{B}$ in $\mathrm{crr}(M)$, let $\{x_1, \ldots, x_k\}$ be the set of variables occurring in $\mathcal{H}$ but not in $\mathcal{B}$. Replace $\mathcal{H} \leftarrow \mathcal{B}$ by the

clause

$$\mathcal{H} \leftarrow \mathcal{B} \wedge \mathsf{dom}(x_1) \wedge \cdots \wedge \mathsf{dom}(x_k).$$

We refer to this clause as the clause *corresponding* to $\mathcal{H} \leftarrow \mathcal{B}$.

(3) Enumerate the Herbrand universe. For each n-ary $f \in \Sigma_f$, add the clauses:

$$\mathsf{dom}(f(x_1, \ldots, x_n)) \leftarrow \mathsf{dom}(x_1) \wedge \cdots \wedge \mathsf{dom}(x_n).$$

The computed set $\mathrm{crr}(M)$ is the *classical range-restricting transformation* of M. It is not difficult to see that $\mathrm{crr}(M)$ is indeed range-restricted for any clause set M. The transformation is sound and complete, that is, M is satisfiable iff $\mathrm{crr}(M)$ is satisfiable [Manthey and Bry, 1988, Bry and Yahya, 2000]. The size of $\mathrm{crr}(M)$ is linear in the size of M and can be computed in linear time.

Perhaps the easiest way to understand the transformation is to imagine we use a BUMG method, for example, hyperresolution. The idea is to build the model(s) during the derivation. The clause added in Step (1) ensures that the domain of interpretation given by the domain predicate dom is non-empty. Step (2) turns clauses into range-restricted clauses. This is done by shielding the variables $\{x_1, \ldots, x_k\}$ in the head, that do not occur negatively, with the added negative domain literals. Clauses that are already range-restricted are unaffected by this step. Step (3) ensures that all elements of the Herbrand universe of the (original) clause set are added to the domain via hyperresolution inference steps.

As a consequence a clause set M with at least one non-nullary function symbols causes hyperresolution derivations to be unbounded for $\mathrm{crr}(M)$, unless M is unsatisfiable. This is a negative aspect of the classical range-restricting transformation. However, the method has been shown to be useful for (domain-)minimal model generation when combined with other techniques [Bry and Yahya, 2000, Bry and Torge, 1998]. In particular, Bry and Torge [1998] use splitting and the δ^*-rule to generate domain minimal models. In the present research we have evaluated the combination of blocking techniques (introduced later in Section 6) with the classical range-restricting transformation crr. This has shown promising empirical results as presented in Section 9.

Next, we introduce a new transformation to range-restricted form. Instead of enumerating the generally infinite Herbrand universe in a bottom-up fashion, the intuition is that it generates terms only as needed.

The transformation involves extracting the non-variable top-level terms in an atom. Let $P(t_1, \ldots, t_n)$ be an atom and suppose $x_1, \ldots, x_n$ are fresh variables. For all $i \in \{1, \ldots, n\}$ let $s_i = t_i$, if t_i is a variable, and $s_i = x_i$, otherwise. The atom $P(s_1, \ldots, s_n)$ is called the *term abstraction of* $P(t_1, \ldots, t_n)$. Let the *abstraction substitution* α be defined by

$$\alpha = \{x_i \mapsto t_i \mid 1 \leq i \leq n \text{ and } t_i \text{ is not a variable}\}.$$

Hence, $P(s_1, \ldots, s_n)\alpha = P(t_1, \ldots, t_n)$, that is, α reverts the term abstraction.

The new *range-restricting transformation*, denoted by rr, of a clause set M is the clause set obtained by carrying out the following steps (explanations and an example are given afterwards):

(0) Initialization. Initially, let $\mathrm{rr}(M) := M$.

(1) Add a constant. Same as Step (1) in the definition of crr.

(2) Domain elements from clause bodies. For each clause $\mathcal{H} \leftarrow \mathcal{B}$ in M and each atom $P(t_1, \ldots, t_n)$ from $\mathcal{B}$, let $P(s_1, \ldots, s_n)$ be the term abstraction of $P(t_1, \ldots, t_n)$ and let α be the corresponding abstraction substitution. Extend $\mathrm{rr}(M)$ by the set

$$\{\mathsf{dom}(x_i)\alpha \leftarrow P(s_1, \ldots, s_n) \mid 1 \leq i \leq n \text{ and } x_i \mapsto t_i \in \alpha\}.$$

(3) Range-restriction. Same as Step (2) in the definition of crr.

(4) Domain elements from Σ_P. For each n-ary P in Σ_p, extend $\mathrm{rr}(M)$ by the set

$$\{\mathsf{dom}(x_i) \leftarrow P(x_1, \ldots, x_n) \mid i \leq i \leq n\}.$$

(5) Domain elements from Σ_f. For each n-ary f in Σ_f, extend $\mathrm{rr}(M)$ by the set

$$\{\mathsf{dom}(x_i) \leftarrow \mathsf{dom}(f(x_1, \ldots, x_n)) \mid 1 \leq i \leq n\}.$$

The intuition of the transformation reveals itself if we think of what happens when using hyperresolution. The idea is again to build model(s) during the derivation, but this time terms are added to the domain only *as necessary*. Steps (1) and (3) are the same as Steps (1) and (2) in the definition of crr. The clauses added in Step (2) cause functional terms that occur negatively in the clauses to be inserted into the domain. Step (4) ensures that positively occurring functional terms are added to the domain, and Step (5) ensures that the domain is closed under subterms.

To illustrate the steps of the transformation consider the following clause.

$$\mathsf{q}(x, \mathsf{g}(x, y)) \vee \mathsf{r}(y, z) \leftarrow \mathsf{p}(\mathsf{a}, \mathsf{f}(x, y), x) \qquad (\dagger)$$

It is added to $\mathrm{rr}(M)$ in Step (0). Suppose the clause added in Step (1) is

$$\mathsf{dom}(a).$$

For Step (2) the term abstraction of the body literal of clause (†) is $\mathsf{p}(x_1, x_2, x)$ and the abstraction substitution is $\alpha = \{x_1 \mapsto \mathsf{a}, x_2 \mapsto \mathsf{f}(x, y)\}$. The clauses added in Step (2) are the following:

$$\begin{aligned} \mathsf{dom}(\mathsf{a}) &\leftarrow \mathsf{p}(x_1, x_2, x) \\ \mathsf{dom}(\mathsf{f}(x, y)) &\leftarrow \mathsf{p}(x_1, x_2, x) \end{aligned} \qquad (\ddagger)$$

Notice that among the four clauses we have so far the clauses (†) and (‡) are not range-restricted. They are however replaced by range-restricted clauses in Step (3), namely:

$$\begin{aligned} \mathsf{q}(x, \mathsf{g}(x, y)) \vee \mathsf{r}(y, z) &\leftarrow \mathsf{p}(\mathsf{a}, \mathsf{f}(x, y), x) \wedge \mathsf{dom}(z) \\ \mathsf{dom}(\mathsf{f}(x, y)) &\leftarrow \mathsf{p}(x_1, x_2, x) \wedge \mathsf{dom}(y). \end{aligned} \qquad (\dagger\dagger)$$

Step (4) generates clauses responsible for inserting the terms that occur in the heads of clauses into the domain. That is, for each $i \in \{1,2,3\}$ and each $j \in \{1,2\}$ these clauses are added.

$$\mathsf{dom}(x_i) \leftarrow \mathsf{p}(x_1, x_2, x_3)$$
$$\mathsf{dom}(x_j) \leftarrow \mathsf{q}(x_1, x_2)$$
$$\mathsf{dom}(x_j) \leftarrow \mathsf{r}(x_1, x_2)$$

For instance, when a model assigns true to the instance $\mathsf{q}(\mathsf{a}, \mathsf{g}(\mathsf{a}, \mathsf{f}(\mathsf{a}, \mathsf{a})))$ of one of the head atoms of the clause (††), then $\mathsf{dom}(\mathsf{a})$ and $\mathsf{dom}(\mathsf{g}(\mathsf{a}, \mathsf{f}(\mathsf{a}, \mathsf{a})))$ are also true. It is not necessary to insert the terms of the instance of the other head atom into the domain. The reason is that it does not matter how these (extra) terms are evaluated, or whether the atom is evaluated to true or false in order to satisfy the disjunction.

The clauses added in Step (4) alone are not sufficient, however. For each term in the domain all its subterms have to be in the domain, too. This is achieved with the clauses obtained in Step (5). That is, for each $j \in \{1,2\}$ these clauses are added.

$$\mathsf{dom}(x_j) \leftarrow \mathsf{dom}(\mathsf{f}(x_1, x_2))$$
$$\mathsf{dom}(x_j) \leftarrow \mathsf{dom}(\mathsf{g}(x_1, x_2))$$

For the purposes of model generation, it is important to note that one particular type of clause in the rr transformation should not be treated as a normal clause. For the equality predicate, Step (4) produces the clauses

$$\mathsf{dom}(x) \leftarrow x \approx y \qquad\qquad \mathsf{dom}(y) \leftarrow x \approx y. \qquad\qquad (\#)$$

Most theorem provers simplify these clauses to $\mathsf{dom}(x)$. As a consequence this can lead to all negative domain literals being resolved away and all clauses containing a positive domain literal to be subsumed. This means range-restriction is undone. This is what happens in SPASS.

Since Step (4) clauses really only need to be added for positively occurring predicate symbols an easy solution involves replacing any positive occurrence of the equality predicate by a predicate symbol $\mathsf{myequal}$ (say), which is fresh in the signature, and adding the clauses

$$\mathsf{dom}(x) \leftarrow \mathsf{myequal}(x, y) \qquad\qquad \mathsf{dom}(y) \leftarrow \mathsf{myequal}(x, y)$$

in Step (4) rather than (#). In addition, the clause set needs to be extended by this definition of $\mathsf{myequal}$.

$$x \approx y \leftarrow \mathsf{myequal}(x, y)$$

This solution has the intended effect of adding terms occurring in positive equality literals to the domain, and prevents other inferences or reductions on $\mathsf{myequal}$. It is not difficult to prove that E-satisfiability is preserved in both directions. We will implicitly use this fact in the proofs below.

Proposition 1 (Completeness of range-restriction)**.** *Let M be any clause set. If $rr(M)$ is satisfiable then M is satisfiable.*

Proof. Suppose $\mathrm{rr}(M)$ is satisfiable. Let I_{rr} be a Herbrand model of $\mathrm{rr}(M)$. We define a quasi-Herbrand interpretation I and show that it is a model of M.

First, the domain of I is defined as the set $|I| = \{t \mid I_{\mathrm{rr}} \models \mathsf{dom}(t)\}$.

Now, to define a total interpretation for the function symbols, we map each n-ary function symbol f in Σ_f to the function $f^I : |I| \times \cdots \times |I| \mapsto |I|$, where, for all $d_1, \ldots, d_n \in |I|$,

$$f^I(d_1, \ldots, d_n) := \begin{cases} f(d_1, \ldots, d_n) & \text{if } f(d_1, \ldots, d_n) \in |I|, \text{ and} \\ c & \text{otherwise.} \end{cases}$$

Here, the constant c is the one mentioned in Step (1) of the transformation. (It is clear that $|I|$ contains c.)

Notice that due to Step (5) the domain $|I|$ must contain for each term all its subterms. An easy consequence is that all terms in $|I|$ are evaluated as themselves, exactly as in Herbrand interpretations. Each other (ground) term is evaluated as some other term from $|I|$. For instance, if $|I| = \{c, f(c)\}$ then $I(f(g(c))) = f^I(I(g(c))) = f^I(g^I(c)) = f^I(c) = f(c)$, since $g(c) \notin |I|$ and by the definition of g^I. We see that f is indeed mapped to a total function over the domain $|I|$, as required.

Regarding the interpretation of the predicate symbols in I, define for every n-ary predicate symbol P in Σ_P and for all $d_1, \ldots, d_n \in |I|$:

$$P(d_1, \ldots, d_n) \in I \text{ iff } P(d_1, \ldots, d_n) \in I_{\mathrm{rr}} \ . \tag{1}$$

That is, the interpretation of the predicate symbols in I is the same as in I_{rr} under the restriction of the domain to $|I| \subseteq |I_{\mathrm{rr}}|$.

It remains to show that I is a model of M. It suffices to pick a clause $\mathcal{H} \leftarrow \mathcal{B}$ from M arbitrarily and show that I satisfies this clause. We do this by assuming that I does not satisfy $\mathcal{H} \leftarrow \mathcal{B}$ and deriving a contradiction.

That I does not satisfy $\mathcal{H} \leftarrow \mathcal{B}$ means there is a valuation μ such that $(I, \mu) \models \mathcal{B}$ but $(I, \mu) \not\models \mathcal{H}$. As usual, a valuation is a (total) mapping from the variables to the domain under consideration.

Because the domain $|I|$ consists of (ground) terms, the valuation μ can be seen as a substitution. Thus, $\mathcal{B}\mu$ is a set of ground atoms, and $\mathcal{B}\mu \subseteq I$ may or may not hold. We show next that if $(I, \mu) \models \mathcal{B}$, as given, then $\mathcal{B}\mu \subseteq I$. In other words, the body is satisfied in I because $|I|$ contains all body atoms $\mathcal{B}\mu$, but not for the reason that I assigns true to some body atom B with some argument term evaluated to c, and that atom being contained in I. An example for the latter case is $|I| = \{\mathsf{c}\}$, $\mathcal{B} = \mathsf{P}(x)$, $I = \{\mathsf{P}(\mathsf{c})\}$ and $\mu = \{x \mapsto \mathsf{a}\}$. Although we have $(I, \mu) \models P(x)$, in essence because $\mathsf{a}^I = \mathsf{c}$, it does *not* hold that $\mathsf{P}(\mathsf{a}) \in I$. The relevance of this result is that it allows syntactically based reasoning further below to show that I is a model of M.

To show $\mathcal{B}\mu \subseteq I$ it suffices to choose any body literal $P(t_1, \ldots, t_n)$ from $\mathcal{B}$ arbitrarily and show $P(t_1, \ldots, t_n)\mu \in I$. Now from $(I, \mu) \models \mathcal{B}$ it follows that $(I, \mu) \models P(t_1, \ldots, t_n)$. Reading μ as a ground substitution this means $P(I(t_1\mu), \ldots, I(t_n\mu)) \in I$. Using the equivalence (1) it follows that $P(I(t_1\mu), \ldots, I(t_n\mu)) \in I_{\mathrm{rr}}$. To show $P(t_1, \ldots, t_n)\mu \in I$, as required, it thus suffices to show $I(t_i\mu) = t_i\mu$, because $P(t_1, \ldots, t_n)\mu \in I$ follows from $P(I(t_1\mu), \ldots, I(t_n\mu)) \in I_{\mathrm{rr}}$ and equivalence (1).

Thus, let us show $I(t_i\mu) = t_i\mu$. By the definition of the interpretation function $\cdot^I$ it is enough to show $t_i\mu \in |I|$ (as said above, terms from $|I|$ are

evaluated to themselves). If t_i is a variable then $t_i\mu \in |I|$ follows from the fact that μ was chosen as a substitution into $|I|$. Assume now that t_i is not a variable and let $P(s_1, \ldots, s_n)$ be the term abstraction of $P(t_1, \ldots, t_n)$ and α its abstraction substitution. By transformation Step (2), $\mathrm{rr}(M)$ includes the clause

$$\mathsf{dom}(x_i)\alpha \leftarrow P(s_1, \ldots, s_n), \tag{2}$$

where $\{x_i \mapsto t_i\} \in \alpha$. By the definition of an abstraction, for all $j \in \{1, \ldots, n\}$, s_j is a fresh variable whenever t_j is not a variable.

Recall from above that $P(I(t_1\mu), \ldots, I(t_n\mu)) \in I_{\mathrm{rr}}$. We are going to show now that with clause (2) this entails $\mathsf{dom}(t_i\mu)$. By the construction of $|I|$ this suffices to prove $t_i\mu \in |I|$, as desired.

Consider the substitution

$$\mu' = \mu\{x_j \mapsto I(t_j\mu) \mid x_j \mapsto t_j \in \alpha\}.$$

It agrees with μ (in particular) when t_j is a variable and otherwise maps the variable x_j to $I(t_j\mu)$.

When t_j is a variable then let $s_j = t_j$ be the definition of an abstraction. This means $s_j\mu' = s_j\mu = t_j\mu = I(t_j\mu)$ (the latter identity holds, again, because μ is a substitution into $|I|$ and elements from $|I|$ evaluate to themselves). When t_j is not a variable then s_j is the variable x_j. By construction of μ' we have $s_j\mu' = x_j\mu' = I(t_j\mu)$. Hence, in both cases $s_j\mu' = I(t_j\mu)$.

Applying the substitution μ' to the clause (2) yields

$$\mathsf{dom}(x_i)\alpha\mu' \leftarrow P(s_1, \ldots, s_n)\mu' \ .$$

With the identities $s_j\mu' = I(t_j\mu)$, the identities $\mathsf{dom}(x_i)\alpha\mu' = \mathsf{dom}(t_i)\mu'$ and the fact that $P(I(t_1\mu), \ldots, I(t_n\mu)) \in I_{\mathrm{rr}}$ it follows that $\mathsf{dom}(t_i)\mu' \in I_{\mathrm{rr}}$. The substitution μ and μ' differ in their domains only on the fresh variables $x_1, \ldots, x_n$. Therefore $\mathsf{dom}(t_i)\mu' = \mathsf{dom}(t_i)\mu$ and $\mathsf{dom}(t_i)\mu \in I_{\mathrm{rr}}$ follows, as desired.

This was the last subgoal to be proven to establish $P(t_1, \ldots, t_n)\mu \in I$, which, in turn, remained to be shown to complete the proof that $\mathcal{B}\mu \subseteq I$.

The next step in the proof is to show that the clause body of the clause in $\mathrm{rr}(M)$ corresponding to $\mathcal{H} \leftarrow \mathcal{B}$ is satisfied by I_{rr}. That clause is the range-restricted version of the clause $\mathcal{H} \leftarrow \mathcal{B}$ in M. According to Step (3) of the transformation it has the form

$$\mathcal{H} \leftarrow \mathcal{B} \wedge \mathsf{dom}(x_1) \wedge \cdots \wedge \mathsf{dom}(x_k) \tag{3}$$

for some variables $x_1, \ldots, x_k$; those occurring in $\mathcal{H}$ but not in $\mathcal{B}$.

From $\mathcal{B}\mu \subseteq I$ as derived above and equivalence (1) it follows that $\mathcal{B}\mu \subseteq I_{\mathrm{rr}}$. Recall that μ is a valuation mapping into the domain $|I|$. Reading it as a substitution gives $x_j\mu \in |I|$, for all $j \in \{1, \ldots, k\}$. From the construction of $|I|$ it follows that $\mathsf{dom}(x_j\mu) \in I_{\mathrm{rr}}$. Together with $\mathcal{B}\mu \subseteq I_{\mathrm{rr}}$ and the fact that I_{rr} is a model of $\mathrm{rr}(M)$, and hence of clause (3), it follows that I_{rr} satisfies $\mathcal{H}\mu$. This means $A\mu \in I_{\mathrm{rr}}$ for some head atom A in $\mathcal{H}$.

The atom A is of the form $Q(s_1, \ldots, s_m)$ for some m-ary predicate symbol Q and terms $s_1, \ldots, s_m$. By Step (4) of the transformation, $\mathrm{rr}(M)$ includes, for all $i \in \{1, \ldots, m\}$ the clause

$$\mathsf{dom}(x_i) \leftarrow Q(x_1, \ldots, x_m). \tag{4}$$

Again by reading μ as a substitution, because I_{rr} is a model of $\mathrm{rr}(M)$, and hence of clause (4), and by the identities $Q(s_1\mu, \ldots, s_m\mu) = Q(s_1, \ldots, s_m)\mu = A\mu \in I_{\mathrm{rr}}$ we conclude $\mathsf{dom}(s_i\mu) \in I_{\mathrm{rr}}$, for all $i \in \{1, \ldots, m\}$. By construction of $|I|$ we have that $s_i\mu \in |I|$. By equivalence (1) it follows that $Q(s_1\mu, \ldots, s_m\mu) \in I$.

Recall that $Q(s_1, \ldots, s_m)$ is a head atom of the clause (3) and hence a head atom of the clause $\mathcal{H} \leftarrow \mathcal{B}$. Further recall that $s_i\mu \in |I|$ entails that $s_i\mu$ is evaluated to itself in I. Together with $Q(s_1\mu, \ldots, s_m\mu) \in I$ this means $(I, \mu) \models Q(s_1, \ldots, s_m)$. This is a contradiction to $(I, \mu) \not\models \mathcal{H}$ as concluded above. The proof is complete. □ □

The proof actually gives a characterization of the models associated with a satisfiable clause set $\mathrm{rr}(M)$.

Corollary 1 (Completeness of range-restriction wrt. E-interpretations)**.** *Let M be any clause set. If $\mathrm{rr}(M)$ is E-satisfiable then M is E-satisfiable.*

Proof. We prove the contra-positive statement. Thus assume M is E-unsatisfiable. Equivalently, $M \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f)$ is unsatisfiable. By Proposition 1, $\mathrm{rr}(M \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f))$ is unsatisfiable. Observe Steps (2) and (3), which are the only ones that apply directly to the given clauses, have no effect on the equality axioms $\mathrm{EAX}(\Sigma_P \cup \Sigma_f)$, except for the reflexivity axiom $x \approx x$, which is replaced by $x \approx x \leftarrow \mathsf{dom}(x)$. The transformed set $\mathrm{rr}(M \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f))$ coincides with

$$\mathrm{rr}(M) \cup (\mathrm{EAX}(\Sigma_P \cup \Sigma_f) \setminus \{x \approx x\}) \cup \{x \approx x \leftarrow \mathsf{dom}(x)\}.$$

Adding back the reflexivity axiom trivially preserves unsatisfiability, that is, with $\mathrm{rr}(M \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f))$ being unsatisfiable, so is

$$\mathrm{rr}(M) \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f) \cup \{x \approx x \leftarrow \mathsf{dom}(x)\}.$$

The clause $x \approx x \leftarrow \mathsf{dom}(x)$ can be deleted because it is subsumed by the clause $x \approx x \in \mathrm{EAX}(\Sigma_P \cup \Sigma_f)$. Hence,

$$\mathrm{rr}(M) \cup \mathrm{EAX}(\Sigma_P \cup \Sigma_f)$$

is unsatisfiable, and so $\mathrm{rr}(M)$ is E-unsatisfiable. □ □

We emphasize that we do not propose to actually use the equality axioms in conjunction with a theorem prover (though they can of course). They serve merely as a theoretical tool to prove completeness of the transformation.

Proposition 2. *Let M be any clause set. Then*

(i) The size of $\mathrm{rr}(M)$ is bounded by a linear function in the size of M.

(ii) $\mathrm{rr}(M)$ can be computed in quadratic time.

(iii) $\mathrm{rr}(M)$ is range-restricted.

By carefully modifying the definition of rr and at the expense of some duplication it is possible to compute the reduction in linear time.

Proposition 2 (iii) confirms that every clauses produced by the rr transformation is range-restricted.

Let us consider another example to get a better understanding of the rr transformation.

$$\mathsf{r}(x) \leftarrow \mathsf{q}(x) \wedge \mathsf{p}(\mathsf{f}(x)). \qquad (*)$$

Applying Steps (2) and (3) of the rr transformation gives us the clause

$$\mathsf{dom}(\mathsf{f}(x)) \leftarrow \mathsf{dom}(x) \wedge \mathsf{p}(y).$$

This clause is splittable into

$$\mathsf{dom}(\mathsf{f}(x)) \leftarrow \mathsf{dom}(x) \quad \text{and} \quad \bot \leftarrow \mathsf{p}(y).$$

The first split component clause is an example of an 'enumerate the Herbrand universe' clause from the crr transformation (Step (2) in the definition of crr). Such clauses are unpleasant because they cause the entire Herbrand universe to be enumerated with BUMG approaches.

Before describing a solution let us analyze the problem further. The main rationale of our rr transformation is to constrain the generation of domain elements and limit the number of inference steps. The general form of clauses produced by Step (2), followed by Step (3), is the following, where $\overline{y} \subseteq \overline{x}$, $\overline{x} \subseteq \overline{y} \cup \overline{z}$ and $\overline{u} \subseteq \overline{z}$.

$$\mathsf{dom}(f(\overline{x})) \leftarrow \mathsf{dom}(y_1) \wedge \ldots \wedge \mathsf{dom}(y_n) \wedge P(\overline{z})$$
$$\mathsf{dom}(f(\overline{u})) \leftarrow P(\overline{z})$$

Clauses of the first form are often splittable (as in the example above), and can produce clauses of the unwanted form

$$\mathsf{dom}(f(\overline{y})) \leftarrow \mathsf{dom}(y_1) \wedge \ldots \wedge \mathsf{dom}(y_n).$$

Suppose therefore that splitting of any clause is forbidden when this splits the negative part of the clause (neither SPASS nor a hypertableaux prover do this anyway). Although the two types of clauses above both *do* reduce the number of terms created, compared to the classical range-restricting transformation, the constraining effect of the first type of clauses is slightly limited. Terms $f(\overline{s})$ are not generated, only when no fact $P(\overline{t})$ is present or has been derived. When a clause $P(\overline{t})$ is present, or as soon as such a clause is derived (for *any* ground terms $\overline{t}$), then terms are freely generated from terms already in the domain with f as the top symbol.

Here is an example of a clause set for which the derivation is infinite on the rr transformation. (The example is an extension of the example above with the clause $\mathsf{p}(\mathsf{b}) \leftarrow \top$.)

$$\mathsf{p}(\mathsf{b}) \leftarrow \top \qquad\qquad \mathsf{r}(x) \leftarrow \mathsf{q}(x) \wedge \mathsf{p}(\mathsf{f}(x))$$

Notice the derivation is infinite on the classical range-restricting transformation as well, due to the generated clauses $\mathsf{dom}(\mathsf{b}) \leftarrow \top$ and $\mathsf{dom}(f(x)) \leftarrow \mathsf{dom}(x)$.

The second type of clauses, $\mathsf{dom}(f(\overline{u})) \leftarrow P(\overline{z})$, are less problematic. Here is a concrete example. For $\bot \leftarrow \mathsf{r}(x, \mathsf{f}(x))$, Step (2) produces the clause

$$\mathsf{dom}(\mathsf{f}(x)) \leftarrow \mathsf{r}(x, y).$$

Although this clause, and the general form, still causes larger terms to be built with hyperresolution type inferences, the constraining effect is larger.

In the next two sections we discuss ways of improving range-restricting transformations further.

5 Shifting Transformation

The clauses introduced in Step (2) of the new rr transformation to range-restricted form use abstraction and insert (possibly a large number of) instantiations of terms occurring in the clause bodies into the domain. These are sometimes unnecessary and can lead to non-termination of BUMG procedures.

The *shifting* transformation introduced next can address this problem. It consists of two sub-transformations, *basic shifting* and *partial flattening*.

If A is an atom $P(t_1, \ldots, t_n)$ then let $\mathsf{not_}A$ denote the atom $\mathsf{not_}P(t_1, \ldots, t_n)$, where $\mathsf{not_}P$ is a fresh predicate symbol which is uniquely associated with the predicate symbol P. If P is the equality symbol $\approx$ we write $\mathsf{not_}P$ as $\not\approx$ and use infix notation.

Now, the *basic shifting transformation* of a clause set M is the clause set $\mathrm{bs}(M)$ obtained from M by carrying out the following steps.

(0) Initialization. Initially, let $\mathrm{bs}(M) := M$.

(1) Shifting deep atoms. Replace each clause in $\mathrm{bs}(M)$ of the form $\mathcal{H} \leftarrow B_1 \wedge \cdots \wedge B_m \wedge \mathcal{B}$, where each atom $B_1, \ldots, B_m$ contains at least one proper functional term and $\mathcal{B}$ contains no proper functional term, by the clause

$$\mathcal{H} \vee \mathsf{not_}B_1 \vee \cdots \vee \mathsf{not_}B_m \leftarrow \mathcal{B}.$$

Each of the atoms $B_1, \ldots, B_m$ is called a *shifted atom*.

(2) Shifted atom consistency. Extend $\mathrm{bs}(M)$ by the clause set

$$\{\bot \leftarrow P(x_1, \ldots, x_n) \wedge \mathsf{not_}P(x_1, \ldots, x_n) \mid P \text{ is the } n\text{-ary predicate symbol of a shifted atom}\}.$$

Notice that we do *not* add clauses complementary to the 'shifted atoms consistency' clauses, that is, $P(x_1, \ldots, x_n) \vee \mathsf{not_}P(x_1, \ldots, x_n) \leftarrow \top$. They could be included but are superfluous.

Let us continue the example given at the end of the previous section. We can use basic shifting to move negative occurrences of functional terms into heads. In the example, clause $(*)$ is replaced by

$$\begin{array}{rl} \mathsf{r}(x) \vee \mathsf{not_p}(\mathsf{f}(x)) \leftarrow \mathsf{q}(x) & \qquad (**) \\ \bot \leftarrow \mathsf{not_p}(x) \wedge \mathsf{p}(x) & \\ \mathsf{dom}(x) \leftarrow \mathsf{r}(x) \qquad\qquad \mathsf{dom}(x) \leftarrow \mathsf{not_p}(x) & \end{array}$$

Even in the presence of an additional clause, say, $\mathsf{q}(x) \leftarrow \top$, which leads to the clauses

$$\mathsf{dom}(\mathsf{a}) \leftarrow \top \qquad\qquad \mathsf{q}(x) \leftarrow \mathsf{dom}(x),$$

termination of BUMG can be achieved.

For instance, in a hyperresolution-like mode of operation and with splitting enabled, the SPASS prover [Weidenbach et al., 2007, 2009] splits the derived clause $\mathsf{r}(\mathsf{a}) \vee \mathsf{not_p}(\mathsf{f}(\mathsf{a}))$, considers the case with the smaller literal $\mathsf{r}(\mathsf{a})$ first

and terminates with a model. This is because a finite completion (model) is found without considering the case of the bigger literal $\mathsf{not_p}(\mathsf{f}(\mathsf{a}))$, which would have added the deeper term $\mathsf{f}(\mathsf{a})$ to the domain. The same behaviour can be achieved, for example, with the KRHyper BUMG prover, a hypertableaux theorem prover [Wernhard, 2003].

As can be seen in the example, the basic shifting transformation trades the generation of new domain elements for a smaller clause body by removing literals from it. Of course, a smaller clause body affects the search space, as then the clause can be used as a premise more often. To (partially) avoid this effect, we propose an additional transformation to be performed prior to the basic shifting transformation.

For a clause set M, the *partial flattening transformation* is the clause set $\mathrm{pf}(M)$ obtained by applying the following steps.

(0) Initialization. Initially, let $\mathrm{pf}(M) := M$.

(1) Reflexivity. Extend $\mathrm{pf}(M)$ by the unit clause $x \approx x \leftarrow \top$.

(2) Partial flattening. For each clause $\mathcal{H} \leftarrow \mathcal{B}$ in $\mathrm{pf}(M)$, let $t_1, \ldots, t_n$ be all top-level terms occurring in the non-equational literals in the body $\mathcal{B}$ that are proper functional terms, for some $n \geq 0$. Let $x_1, ..., x_n$ be fresh variables. Replace the clause $\mathcal{H} \leftarrow \mathcal{B}[t_1, \ldots, t_n]$ by the clause

$$\mathcal{H} \leftarrow \mathcal{B}[x_1, \ldots, x_n] \wedge t_1 \approx x_1 \wedge \cdots \wedge t_n \approx x_n.$$

It should be noted that the equality symbol $\approx$ need not be interpreted as equality, but could. (Un-)satisfiability (and logical equivalence) is preserved even when reading it just as 'unifiability'. This can be achieved by the clause $x \approx x \leftarrow \top$. One should however note that the reflexivity clause is not compatible with introducing the $\mathsf{myequal}$ predicate, so this might not always be a possibility. (In our implementation, for this reason the reflexivity clause is not added.)

In our running example, applying the transformations pf, bs and rr, in this order, yields the following clauses (among other clauses, which are omitted because they are not relevant to the current discussion).

$$\mathsf{r}(x) \vee \mathsf{f}(x) \not\approx u \leftarrow \mathsf{q}(x) \wedge \mathsf{p}(u) \qquad \mathsf{dom}(x) \leftarrow x \not\approx y \qquad \mathsf{dom}(x) \leftarrow \mathsf{r}(x)$$
$$\bot \leftarrow x \not\approx y \wedge x \approx y \qquad \mathsf{dom}(y) \leftarrow x \not\approx y$$

Observe that the first clause is more restricted than the clause $(**)$ above because of the additional body literal $\mathsf{p}(u)$.

The reason for not extracting constants during partial flattening is that adding them to the domain does not cause non-termination of BUMG methods. It is preferable to leave them in place in the body literals because they have a stronger constraining effect than the variables introduced otherwise.

Extracting top-level terms from equations has no effect at all. Consider the unit clause $\bot \leftarrow f(a) \approx b$, and its partial flattening $\bot \leftarrow x \approx b \wedge f(a) \approx x$. Applying basic shifting yields $f(a) \not\approx x \leftarrow x \approx b$, and, hyperresolution with $x \approx x \leftarrow \top$ gives $f(a) \not\approx b \leftarrow \top$. This is the same result as obtained by the transformations as defined. This explains why top-level terms of equational literals are excluded from the definition. (One could consider using 'standard' flattening, that is, recursively extracting terms, but this does not lead to any improvements over the defined transformations.)

Finally, we combine basic shifting and partial flattening to give the *shifting transformation*, formally defined by $\mathrm{sh} := \mathrm{pf} \circ \mathrm{bs}$, that is, $\mathrm{sh}(M) = \mathrm{bs}(\mathrm{pf}(M))$, for any clause set M.

Proposition 3 (Completeness of shifting)**.** *Let M be any clause set. If $sh(M)$ is satisfiable then M is satisfiable.*

Proof. Not difficult, since bs (basic shifting) can be seen to be a structural transformation and pf (partial flatting) is a form of term abstraction. □ □

Corollary 2 (Completeness of shifting wrt. E-interpretations)**.** *Let M be any clause set. If $sh(M)$ is E-satisfiable then M is E-satisfiable.*

Proof. Using the same line of argument as in the proof of Corollary 1, proving preservation of E-satisfiability can be reduced to proving preservation of satisfiability by means of the equality axioms (observe that the shifting transformation does not modify the equality axioms). □ □

6 Blocking

The final transformation introduced in this paper is called *blocking* and provides a mechanism for detecting recurrence in the derived models. The blocking transformation is designed to realize a 'loop check' for the construction of a domain, by capitalizing on available, powerful equality reasoning technology and redundancy criteria from saturation-based theorem proving. To be suitable, a resolution-based prover, for instance, should support *hyperresolution-style inference, strong equality inference (for example, superposition or ordered rewriting), splitting, and the possibility to search for split-off equations first and standard redundancy elimination techniques.*

The basic idea behind blocking is to add clauses that cause a case analysis of the form $s \approx t$ versus $s \not\approx t$, for (ground) terms s and t. Although such a case analysis obviously leads to a bigger search space, it provides a powerful technique to detect finite models with a BUMG prover. This is because in the case that $s \approx t$ is assumed, this new equation may lead to rewriting of otherwise infinitely many terms into one single term. To make this possible, the prover must support the above features, including notably splitting. Among resolution theorem provers splitting has become standard. Splitting was first available in the saturation-based prover SPASS [Weidenbach et al., 2007, 2009], but is now also part of VAMPIRE [Riazanov and Voronkov, 2002] and E [Schulz, 2013]. Splitting is an integral part of the hypertableau prover E-KRHyper [Baumgartner et al., 2007, Pelzer and Wernhard, 2007].

Blocking has the same goal as the *unsound theorem proving* technique introduced first in Lynch [2004]. Instances of unsound theorem proving exemplified in Lynch [2004] include replacing a clause by one that subsumes it, and by adding equations for joining equivalence classes in the abstract congruence closure framework. Unsound theorem proving has been incorporated later in DPLLT-based theorem proving Bonacina et al. [2011].

In the following we introduce four different, but closely related, blocking transformations, called *subterm domain blocking*, *subterm predicate blocking*, *unrestricted domain blocking* and *unrestricted predicate blocking*. Subterm domain

blocking was introduced in the short version of this paper under the name blocking [Baumgartner and Schmidt, 2006]. Subterm predicate blocking is inspired by and related to the blocking technique described in Hustadt and Schmidt [1999]. Unrestricted domain blocking is the first-order version of the unrestricted blocking rule introduced in Schmidt and Tishkovsky [2007] and used for developing terminating tableau calculi for logics with the effective finite model property in Schmidt and Tishkovsky [2008, 2011].

6.1 Subterm Domain Blocking

By definition, the *subterm domain blocking transformation* of a clause set M is the clause set $\mathrm{sdb}(M)$ obtained from M by carrying out the following steps.

(0) Initialization. Initially, let $\mathrm{sdb}(M) := M$.

(1) Axioms describing the subterm relationship. Let sub be a 'fresh' binary predicate symbol not in Σ_P. Extend $\mathrm{sdb}(M)$ by

$$\mathsf{sub}(x, x) \leftarrow \mathsf{dom}(x)$$

and, for every n-ary function symbol $f \in \Sigma_f$ and all $i \in \{1, \ldots, n\}$, add the clauses

$$\mathsf{sub}(x, f(x_1, \ldots, x_n)) \leftarrow \mathsf{sub}(x, x_i) \wedge \mathsf{dom}(x) \wedge \mathsf{dom}(f(x_1, \ldots, x_n)).$$

(2) Subterm equality case analysis. Extend $\mathrm{sdb}(M)$ by these clauses.

$$\begin{aligned} x \approx y \vee x \not\approx y &\leftarrow \mathsf{sub}(x, y) \\ &\leftarrow x \approx y \wedge x \not\approx y \end{aligned}$$

The subterm domain blocking transformation allows to contemplate whether two domain elements that are in a subterm relationship should be identified and merged, or not.

This blocking transformation preserves range-restrictedness. In fact, because the dom predicate symbol is mentioned in the definition, the blocking transformation can be applied meaningfully only in combination with range-restricting transformations.

Reading $\mathsf{sub}(s, t)$ as 's is a subterm of t', Step (1) in the blocking transformation might seem overly involved, because an apparently simpler specification of the subterm relationship for the terms of the signature Σ_f can be given. Namely:

$$\mathsf{sub}(x, x) \leftarrow \mathsf{dom}(x) \qquad \mathsf{sub}(x, f(x_1, x_2 \ldots, x_n)) \leftarrow \mathsf{sub}(x, x_i)$$

for every n-ary function symbol $f \in \Sigma_f$ and all $i \in \{1, \ldots, n\}$. This clause set is range-restricted. Yet, this specification is not suitable for our purposes. The problem is that the second clause introduces proper functional terms.

For example, for a given constant a and a unary function symbol f, when just $\mathsf{dom}(\mathsf{a})$ alone has been derived, a BUMG procedure derives an infinite sequence of clauses:

$$\mathsf{sub}(\mathsf{a}, \mathsf{a}), \mathsf{sub}(\mathsf{a}, \mathsf{f}(\mathsf{a})), \mathsf{sub}(\mathsf{a}, \mathsf{f}(\mathsf{f}(\mathsf{a}))), \ldots.$$

This does not happen with the specification in Step (1). It ensures that conclusions of BUMG inferences involving sub are about terms currently in the domain, and the domain is always finite.

To justify the clauses added in Step (2) we continue this example and suppose an interpretation that contains $\mathsf{dom}(\mathsf{a})$ and $\mathsf{dom}(\mathsf{f}(\mathsf{a}))$. These might have been derived earlier in the run of a BUMG prover. Then, from the clauses added by blocking, the (necessarily ground) disjunction

$$\mathsf{f}(\mathsf{a}) \approx \mathsf{a} \vee \mathsf{f}(\mathsf{a}) \not\approx \mathsf{a} \leftarrow \top$$

is derivable.

Now, it is important to use a BUMG prover with support for splitting and to equip it with an appropriate search strategy. In particular, when deriving a disjunction such as the one above, the $\approx$-literal should be split off and the clause set obtained in this case should be searched *first*. The reason is that the (ground) equation $\mathsf{f}(\mathsf{a}) \approx \mathsf{a}$ thereby obtained can then be used for simplification and redundancy testing purposes. For example, should $\mathsf{dom}(\mathsf{f}(\mathsf{f}(\mathsf{a})))$ be derivable now (in the current branch), then any prover based a modern, saturation-based theory of equality reasoning is able to prove it redundant from $\mathsf{f}(\mathsf{a}) \approx \mathsf{a}$ and $\mathsf{dom}(\mathsf{a})$. Consequently, the domain is *not* be extended *explicitly*. The information that $\mathsf{dom}(\mathsf{f}(\mathsf{f}(\mathsf{a})))$ is in the domain is however implicit via the theory of equality.

6.2 Subterm Predicate Blocking

Subterm domain blocking defined in the previous section applies blocking to *domain terms* where one is a proper subterm of the other. The idea of the *subterm (unary) predicate blocking transformation* is similar, but it merges only the (sub)terms in the extension of *unary* predicate symbols different to dom in the current interpretation.

Subterm predicate blocking is defined as follows:

(0) Initialization. Initially, let $\mathrm{spb}(M) := M$.

(1) Axioms describing the subterm relationship. Same as Step (1) in the definition of sdb.

(2) Subterm equality case analysis. Extend $\mathrm{spb}(M)$ by these clauses, for each unary predicate symbol $p \in \Sigma_P$. (Recall that Σ_P does not contain dom.)

$$x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y) \wedge p(x) \wedge p(y)$$

Finally, add the clause

$$\leftarrow x \approx y \wedge x \not\approx y$$

to $\mathrm{spb}(M)$.

Observe that the only difference between this transformation and the subterm domain blocking transformation lies in Step (2). The clauses $x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y) \wedge p(x) \wedge p(y)$ added here are obviously more restrictive than their

counterpart $x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y)$ in the definition of the subterm domain blocking transformation sdb.

That subterm predicate blocking is strictly more restrictive can be seen from the following example, which also helps to explain the rationale behind this transformation.

$$\mathsf{p}(\mathsf{a}) \leftarrow \qquad\qquad \mathsf{q}(\mathsf{f}(x)) \leftarrow \mathsf{p}(x)$$

Any BUMG prover terminates on the transformed clause set and returns the model

$$\{\mathsf{dom}(\mathsf{a}),\ \mathsf{dom}(\mathsf{f}(\mathsf{a})),\ \mathsf{p}(\mathsf{a}),\ \mathsf{q}(\mathsf{f}(\mathsf{a})),\ \mathsf{sub}(\mathsf{a}, \mathsf{f}(\mathsf{a}))\}.$$

Notice that the subterm predicate blocking transformation includes the clauses

$$x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y) \wedge \mathsf{p}(x) \wedge \mathsf{p}(y)$$
$$x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y) \wedge \mathsf{q}(x) \wedge \mathsf{q}(y).$$

These are however only applicable for $\mathsf{sub}(\mathsf{s}, \mathsf{s})$, $\mathsf{p}(s)$ and $\mathsf{q}(s)$ which only lead to redundant BUMG inferences. The motivation behind these clauses is to block two p-literals (say) only when there are two literals $\mathsf{p}(s)$ and $\mathsf{p}(t)$ where s is a subterm of t. Conversely, if no such loop comes up, as in the example above, there is no reason for blocking. By contrast, the subterm domain blocking transformation sdb with its clause $x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y)$ would be applicable even for distinct terms, leading to the (unnecessary) split into the cases $\mathsf{a} \approx \mathsf{f}(\mathsf{a})$ and $\mathsf{a} \not\approx \mathsf{f}(\mathsf{a})$.

From a more general perspective, the spb transformation is motivated by the application to description logic knowledge bases [Hustadt and Schmidt, 1999, Baader and Sattler, 2001]. Often, such knowledge bases do not contain cyclic definitions, or only few definitions are cyclic. The subterm predicate transformation aims to apply blocking only to concepts (unary predicates) with cyclic definitions. Below, in Section 6.5, we discuss a description logic example to highlight the differences between the various blocking transformations.

6.3 Unrestricted Domain Blocking

The two previous 'subterm' variants of the blocking transformation allow to speculatively identify terms *and their subterms*. The 'unrestricted' variants introduced next differ from both by allowing speculative identifications of *any* two terms.

For the 'domain' variant, called *unrestricted domain blocking transformation*, the definition is as follows.

(0) Initialization. Initially, let $\mathrm{udb}(M) := M$.

(1) Domain elements equality case analysis. Extend $\mathrm{udb}(M)$ by these clauses.

$$x \approx y \vee x \not\approx y \leftarrow \mathsf{dom}(x) \wedge \mathsf{dom}(y)$$
$$\leftarrow x \approx y \wedge x \not\approx y$$

There is a clear trade-off between this transformation and the subterm domain blocking transformation sdb. On the one hand, the unrestricted domain

blocking transformation induces a larger search space, as the bodies of the clauses $x \approx y \vee x \not\approx y \leftarrow \mathsf{dom}(x) \wedge \mathsf{dom}(y)$ are less constrained than their counterparts in the subterm domain blocking transformation. This becomes obvious after extending the clause body of $x \approx y \vee x \not\approx y \leftarrow \mathsf{sub}(x, y)$ from the sdb transformation with $\mathsf{dom}(x) \wedge \mathsf{dom}(y)$, which does not change anything. On the other hand, the unrestricted domain blocking transformation enables the finding of models with smaller domains. This means fewer congruence classes on the Herbrand terms are induced by the equality relation $\approx$. As our experiments show, such models can often be found quicker in satisfiable problems, even for the crr transformation.

Using the ideas of the termination proof in Schmidt and Tishkovsky [2013] for semantic ground tableau with unrestricted domain blocking for description logics with the expressive power similar to the two-variable fragment of first-order logic, it can be shown BUMG with unrestricted domain blocking can return finite models, if they exist, even for problems of undecidable fragments. Carrying over also the results in Schmidt and Tishkovsky [2008] implies unrestricted domain blocking can be used in BUMG methods to return domain minimal models for logics with the effective finite model property.

6.4 Unrestricted Predicate Blocking

The definition of the last variant of blocking, the *unrestricted (unary) predicate blocking transformation*, is as follows.

(0) Initialization. Initially, let $\mathrm{upb}(M) := M$.

(1) Term equality case analysis. Extend $\mathrm{upb}(M)$ by these clauses, for each unary predicate symbol $p \in \Sigma_P$.

$$x \approx y \vee x \not\approx y \leftarrow p(x) \wedge p(y)$$

Finally, add the clause

$$\leftarrow x \approx y \wedge x \not\approx y$$

to $\mathrm{upb}(M)$.

This transformation allows to equate any two (distinct) terms in a p-relation, if there are any. The motivation is a combination of the above, to block cycles on p-literals if they arise, and to compute models with small domains.

6.5 Comparison on an Example

It is instructive to compare the effects on the returned models of the four blocking transformations on an example from description logics. To this end, consider the description logic knowledge base (left) and its translation into clause logic (right) in Table 1. Notice that the cycle in the inclusion statements in the TBox (for p_1 and p_2) means some form of blocking is needed for decidability in tableau-based description logic systems. Likewise, blocking is needed to force BUMG methods to terminate on the translated clause form. Any of the four blocking transformations defined above suffice. Table 2 summarizes the behaviour of these transformations, in terms of interesting relations in the computed model.

Table 1: Sample description logic knowledge base and clausal form

TBox	ABox
$\mathsf{p}_1 \sqsubseteq \exists\, \mathsf{r}.\mathsf{p}_2$	$\mathsf{p}_1(\mathsf{a})$
$\mathsf{p}_2 \sqsubseteq \exists\, \mathsf{r}.\mathsf{p}_1$	$\mathsf{p}_1(\mathsf{b})$
$\mathsf{p}_1 \sqsubseteq \exists\, \mathsf{s}.\mathsf{q}$	

$$\mathsf{p}_2(\mathsf{f}(x)) \leftarrow \mathsf{p}_1(x) \qquad \mathsf{q}(\mathsf{h}(x)) \leftarrow \mathsf{p}_1(x)$$
$$\mathsf{r}(x, \mathsf{f}(x)) \leftarrow \mathsf{p}_1(x) \qquad \mathsf{s}(x, \mathsf{h}(x)) \leftarrow \mathsf{p}_1(x)$$
$$\mathsf{p}_1(\mathsf{g}(x)) \leftarrow \mathsf{p}_2(x) \qquad \mathsf{p}_1(\mathsf{a}) \leftarrow$$
$$\mathsf{r}(x, \mathsf{g}(x)) \leftarrow \mathsf{p}_2(x) \qquad \mathsf{p}_1(\mathsf{b}) \leftarrow$$

Table 2: Partial truth assignments in models computed for the sample knowledge base

Blocking	dom	$\approx$	p_1	p_2	q
sdb	a, b	$\mathsf{f}(\mathsf{a}) \approx \mathsf{a}, \mathsf{f}(\mathsf{b}) \approx \mathsf{b},$ $\mathsf{g}(\mathsf{a}) \approx \mathsf{a}, \mathsf{g}(\mathsf{b}) \approx \mathsf{b},$ $\mathsf{h}(\mathsf{a}) \approx \mathsf{a}, \mathsf{h}(\mathsf{b}) \approx \mathsf{b}$	a, b	a, b	a, b
spb	$\mathsf{a}, \mathsf{b},$ $\mathsf{f}(\mathsf{a}), \mathsf{f}(\mathsf{b}),$ $\mathsf{h}(\mathsf{a}), \mathsf{h}(\mathsf{b})$	$\mathsf{g}(\mathsf{f}(\mathsf{a})) \approx \mathsf{a},$ $\mathsf{g}(\mathsf{f}(\mathsf{b})) \approx \mathsf{b}$	a, b	$\mathsf{f}(\mathsf{a}), \mathsf{f}(\mathsf{b})$	$\mathsf{h}(\mathsf{a}), \mathsf{h}(\mathsf{b})$
udb	b	$\mathsf{a} \approx \mathsf{b}, \mathsf{f}(\mathsf{b}) \approx \mathsf{b},$ $\mathsf{g}(\mathsf{b}) \approx \mathsf{b}, \mathsf{h}(\mathsf{b}) \approx \mathsf{b}$	b	b	b
upb	$\mathsf{b}, \mathsf{f}(\mathsf{b}), \mathsf{h}(\mathsf{b})$	$\mathsf{a} \approx \mathsf{b},$ $\mathsf{g}(\mathsf{f}(\mathsf{b})) \approx \mathsf{b}$	b	$\mathsf{f}(\mathsf{b})$	$\mathsf{h}(\mathsf{b})$

When comparing in detail the blocking techniques developed for description logics it becomes clear that

$$\text{the transformations } \mathrm{rr} \circ \tau \text{ and } \mathrm{sh} \circ \mathrm{rr} \circ \tau, \text{ for } \tau \in \{\mathrm{sdb}, \mathrm{spb}, \mathrm{udb}, \mathrm{upb}\},$$

when applied to a knowledge base with the finite model property, in conjunction with a suitable BUMG method (see above), can be refined to simulate various forms of standard blocking techniques used in description logic systems, including subset ancestor blocking and equality ancestor blocking, cf. Hustadt and Schmidt [1999], Schmidt and Tishkovsky [2013] and Khodadadi et al. [2013]. Because standard loop checking mechanisms used in description logic systems do not require backtracking, appropriate search strategies and restrictions for performing inferences and applying blocking need to be used.

An advantage of our approach to blocking as opposed to blocking without equality reasoning used in mainstream description logic systems [Baader and Sattler, 2001] is that it applies to any first-order clause set, not only to clauses from the translation of description logic problems. This makes the approach very general and widely applicable.

For instance, our approach makes it possible to extend description logics with arbitrary (first-order expressible) 'rule' languages. 'Rules' provide a connection to (deductive) databases and are being used to represent information that is currently not expressible in the description logics associated with OWL DL. The specification of many natural properties of binary relations and complex statements involving binary relations are outside the scope of most current description logic systems. An example is the statement: individuals who live and

work at the same location are home workers. This can be expressed as a Horn rule (clause)

$$\mathsf{homeWorker}(x) \leftarrow \mathsf{work}(x, y) \wedge \mathsf{live}(x, z) \wedge \mathsf{loc}(y, w) \wedge \mathsf{loc}(z, w),$$

but, with some exceptions [Hustadt et al., 1999, Weidenbach et al., 2007], is not expressible in current description logic systems.

7 Soundness and Completeness of the Transformations

Each of the blocking transformations is complete:

Proposition 4 (Completeness of blocking wrt. E-interpretations)**.** *Let M be any clause set. For all $\tau \in \{sdb,\ spb,\ udb,\ upb\}$, if $\tau(M)$ is E-satisfiable then M is E-satisfiable.*

Proof. Not difficult, as $M \subseteq \tau(M)$ by definition. □ □

The converse, that is, soundness of the transformation, is easy to prove. One basically needs to observe that the clauses added in respectively Steps (2) and (1) of the blocking transformations, realize a case distinction over whether two terms are equal or not. Trivially, one of the two cases always holds.

Putting all the transformations and the corresponding results together we can state the main theoretical result of the paper.

Theorem 1 (Completeness of the combined transformations with respect to E-interpretations)**.** *Let M be a clause set and suppose tr is any of the transformations in $\{rr, sh \circ rr\} \cup \{rr \circ \tau, sh \circ rr \circ \tau \mid \tau \in \{sdb,\ spb,\ udb,\ upb\}\}$ or $\{crr, sh \circ crr\} \cup \{crr \circ \tau, sh \circ crr \circ \tau \mid \tau \in \{sdb,\ spb,\ udb,\ upb\}\}$. Then:*

(i) $tr(M)$ is range-restricted.

(ii) $tr(M)$ can be computed in quadratic time.

(iii) If $tr(M) \cup \{x \approx x \leftarrow \mathsf{dom}(x)\}$ is E-satisfiable then M is E-satisfiable.

The reverse directions of (iii), that is, soundness of the respective transformations, hold as well. The proofs are either easy or completely standard.

By carefully modifying the definition of rr it is possible to compute the reductions in linear time.

Proposition 5. *Let M and tr be as in the previous result. Then:*

(i) The size of $tr(M)$ is bounded by a linear function in the size of M.

(ii) $tr(M)$ can be computed in linear time.

8 Decidability of BS classes

The Bernays-Schönfinkel class can be decided using transformations into range-restricted clauses. Formulae in the Bernays-Schönfinkel class are conjunctions of function-free and equality-free formulae of the form $\exists^*\forall^*\psi$, where ψ is free of quantifiers. A clause is a *BS clause* iff all functional terms occurring in it are constants.

It is proved in Schmidt and Hustadt [2005] that hyperresolution and any refinements decide the class of range-restricted BS clauses without equality. Here assume that the language includes equality.

Theorem 2. *The class of range-restricted BS clauses (with equality), is decidable by hyperresolution (and paramodulation) and all refinements.*

This means all refinements of hyperresolution (and some form of equality reasoning) combined with any translation into range-restricted clauses is a decision procedure for the BS class.

Therefore:

Corollary 3. *Let M be any set of BS clauses, and suppose tr is any of the transformations in $\{rr, sh \circ rr\} \cup \{rr \circ \tau, sh \circ rr \circ \tau \mid \tau \in \{sdb,\ spb,\ udb,\ upb\}\}$ and $\{crr, sh \circ crr\} \cup \{crr \circ \tau, sh \circ rr \circ \tau \mid \tau \in \{sdb,\ spb,\ udb,\ upb\}\}$. Then:*

(i) Hyperresolution and all refinements decide $tr(M)$.

(ii) All BUMG methods decide M.

Since there are linear transformations of first-order formulae into clausal form, and since all the tr transformations are effective reductions of first-order clauses into range-restricted clauses, we obtain the following result.

Theorem 3.

(i) There is a quadratic (linear), satisfiability equivalence preserving transformation of any formula in the Bernays-Schönfinkel class, and any set of BS clauses, into a set of range-restricted BS clauses.

(ii) All procedures based on hyperresolution or BUMG decide the class of BS formulae and the class of BS clauses.

In Schmidt and Hustadt [2005] a similar but different transformation is used to prove this result for hyperresolution and BS without equality. In fact, what is crucial for deciding the BS class is a grounding method. This can be achieved by any form of range-restriction and hyperresolution-like inferences. Theorem 3.(ii) can therefore be strengthened to include also any instantiation-based method, in particular also methods using on-the-fly instantiation such as semantic Smullyan-type tableaux.

9 Experimental Evaluation

We have implemented the transformations described in the previous sections and carried out experiments on problems from the TPTP library, Version 6.0.0.

The implementation, in SWI-Prolog, is called Yarralumla (Yet another range-restriction avoiding loops under much less assumptions). Since the transformations introduced in this paper are defined for clausal problems we have selected for the experiments all the CNF problems from the TPTP suite.

In our initial research [Baumgartner and Schmidt, 2006] we used Yarralumla with the MSPASS theorem prover, Version 2.0g.1.4 [Hustadt and Schmidt, 2000]. As the extra features of MSPASS have in the mean time been integrated into the SPASS theorem prover [Weidenbach et al., 2007] and SPASS has significantly evolved since Version 2.0, for the present paper we combined Yarralumla with SPASS Version 3.8d as a BUMG system.

For that purpose we modified the code of SPASS in a number of ways. We added one new flag to activate splitting on positive ground equality literals in positive non-Horn clauses. The main inference loop was updapted so that finding a splitting clause and applying splitting has highest priority (unchanged) followed immediately by picking a non-positive blocking clause, that is, clauses of the form $s \approx t \vee H_1 \vee \cdots \vee H_m \leftarrow B_1 \wedge \cdots \wedge B_k$ for $m \geq 0$ and $k > 0$, and performing inferences with it. The selection of splitting clauses was adapted so that positive blocking clauses are always selected, when there are any. Moreover, the first equality literal is split upon. Positive blocking clauses are ground clauses of the form $s \approx t \vee H_1 \vee \cdots \vee H_m$, where $m \geq 1$. This adaptation ensures blocking is performed eagerly to keep the set of ground terms small. The tests with Yarralumla were performed using ordered resolution and superposition with selection of at least one negative literal, forward and backward rewriting, unlimited splitting and matching replacement resolution, subsumption deletion and various other simplification rules. This means the inferences are performed in an ordered hyperresolution-style with eager splitting and forward and backward ground rewriting. The derivations constructed are thus BUMG tree derivations, the proofs produced are BUMG refutation proofs, and the models returned are BUMG models.

We also tested SPASS Version 3.8d in auto mode on the sample. In auto mode SPASS used ordered resolution with dynamic selection. SPASS automatically turned off splitting for non-Horn clauses. Dynamic selection means typically literals were only selected if multiple maximal literals occur in a clause. This means the behaviour of SPASS in auto mode was very different to that of SPASS-Yarralumla, which always selected a literal in clauses with non-empty negative part. The changes to SPASS in SPASS-Yarralumla meant that splitting was performed eagerly and blocking clauses were targeted, which was not the case with SPASS in auto mode. We tested SPASS in auto mode only on the original files (translated from TPTP syntax to SPASS syntax).

The experiments were run on a cluster of 128 Dell PowerEdge M610 Blade Servers each with two Intel Xeon E5620 2.4 GHz processors and 48 GiB main memory each. The time limit was ten minutes (CPU time).

SPASS-Yarralumla can be downloaded from
`http://www.cs.man.ac.uk/~schmidt/spass-yarralumla/`.

9.1 Results

Tables 3 and 4 summarize the results for satisfiable clausal problems in the TPTP library, measuring the number of problems solved with in the time limit. The columns with the heading '#' give the number of problems in the TPTP

Table 3: Number of problems solved on satisfiable problems, by TPTP categories.

Category	#	rr	rr ∘ sdb	rr ∘ udb	rr ∘ spb	rr ∘ upb	sh ∘ rr	sh ∘ rr ∘ sdb	sh ∘ rr ∘ udb	sh ∘ rr ∘ spb	sh ∘ rr ∘ upb	crr	crr ∘ sdb	crr ∘ udb	crr ∘ spb	crr ∘ upb	sh ∘ crr	sh ∘ crr ∘ sdb	sh ∘ crr ∘ udb	sh ∘ crr ∘ spb	sh ∘ crr ∘ upb	auto
ALG	37	-	-	-	-	-	-	-	-	-	-	-	**1**	**1**	-	-	-	-	-	-	-	**3**
ANA	2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
BOO	17	-	**5**	**5**	-	-	-	4	**5**	-	-	-	3	**6**	-	-	-	3	**6**	-	-	2
CAT	10	2	**6**	4	4	2	2	**4**	**4**	**4**	2	2	3	**4**	2	2	2	3	**4**	3	2	2
COL	6	-	-	-	-	-	-	-	-	-	-	-	1	**3**	-	-	-	1	**3**	-	-	-
GEO	17	-	-	-	-	-	-	-	**1**	-	-	-	**1**	-	-	-	-	-	-	-	-	-
GRP	85	49	55	**60**	49	50	50	53	**57**	50	51	49	55	**61**	49	50	49	52	**56**	49	50	-
HEN	3	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
HWC	2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
HWV	39	2	2	2	2	2	2	2	2	2	2	2	**3**	**3**	2	2	2	2	**3**	2	2	-
KRS	13	1	**8**	**8**	7	5	7	**8**	**8**	**8**	**8**	-	7	**8**	7	5	-	7	7	7	5	**8**
LAT	62	-	**10**	4	-	-	-	**6**	5	-	-	-	25	**32**	-	-	-	24	**30**	-	-	-
LCL	44	1	**4**	1	2	1	2	**4**	3	2	2	-	2	**5**	-	-	1	2	**5**	1	1	-
LDA	26	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MGT	11	1	6	**9**	6	1	4	7	7	**8**	4	-	**7**	5	5	-	-	4	**5**	4	-	8
MSC	2	1	1	1	1	1	1	1	1	1	1	-	1	1	1	1	-	1	1	1	1	1
NLP	236	52	130	**137**	128	63	99	191	**195**	188	110	22	**125**	**125**	114	33	22	118	**125**	104	33	**198**
NUM	8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
PLA	6	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
PUZ	27	7	**8**	**8**	7	7	7	**8**	**8**	7	7	7	**8**	**8**	7	7	7	**8**	**8**	7	7	7
REL	1	-	**1**	**1**	-	-	-	**1**	**1**	-	-	-	**1**	**1**	-	-	-	**1**	**1**	-	-	-
RNG	14	1	1	1	1	1	1	1	1	1	1	1	1	**2**	1	1	1	**3**	1	1	1	-
ROB	5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
SCT	3	2	**3**	2	**3**	**3**	2	2	**3**	2	**3**	2	**3**	**3**	**3**	**3**	2	2	**3**	2	2	**3**
SET	12	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
SWC	1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
SWV	147	2	6	**8**	7	2	2	2	**7**	2	2	-	5	**7**	5	-	-	4	**6**	4	-	**79**
SWW	39	-	7	**8**	-	-	-	**2**	1	1	-	-	-	**4**	-	-	-	-	-	-	-	-
SYN	223	54	132	**147**	126	117	56	132	**145**	125	117	54	134	**148**	126	110	55	131	**143**	124	111	99
SYO	9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
TOP	19	-	2	**3**	1	-	-	**1**	**1**	-	-	-	**2**	**2**	1	-	-	1	**3**	-	-	**6**
Total	1126	176	388	**410**	345	256	236	430	**456**	402	311	140	389	**430**	324	215	142	368	**411**	310	216	417

categories and the different TPTP rating ranges. The subsequent columns give the number of problems solved within the time limit. The results are presented for the different BUMG methods that were used. For example, sh∘rr∘sdb refers to the method based on the transformation defined by the new range-restriction transformation, shifting and subterm domain blocking. To evaluate the effect of the different forms of blocking the results are grouped into groups of five: no blocking, subterm domain blocking (sdb), unrestricted domain blocking (udb), subterm predicate blocking (spb) and unrestricted predicate blocking (spb). In each group the first column provides the *baseline* for that group. The last column with the heading 'auto' gives the results for runs of SPASS Version 3.8d in auto mode on the original input files. The runtimes for the problems solved spanned the whole range, from less than one second to all of the time allowed.

The best results in each group in each row are highlighted in bold font. The underlined values are the best results for all methods including SPASS in auto mode. As expected the worst results in each group were obtained for the baseline transformations without blocking. This confirms the expectation that blocking is an essential technique for BUMG methods. Among the different blocking techniques the best results were obtained with unrestricted domain blocking in all four groups. Overall, the best result was obtained for the combination with rr and shifting, i.e., sh ∘ rr ∘ udb, solving 6.0% more problems than the second best method, crr∘udb using the classical range-restriction transformation without shifting, and nearly 11% more problems than the transformations rr ∘ udb and sh ∘ crr ∘ udb. This means shifting had a significant positive effect in combination with the new range-restriction transformation, but less so in combination with classical range-restriction. The positive effect of shifting could also be seen for the number of problems solved without blocking for rr and sh∘rr (34% improvement).

The good results for crr ∘ udb show the value of classical range-restriction. In the LAT category, crr ∘ udb solved 32 problems, whereas sh ∘ rr ∘ udb solved only 5 problems. This seems to indicate there was a trade-off between using the crr transformation and the rr transformation in combination with shifting, but also showed the virtues of unrestricted domain blocking as a universal technique for BUMG. SPASS in auto mode fared very well in the SWV category, where 79 problems were solved compared to 7–8 problems for the best BUMG methods. Overall SPASS in auto mode solved 9% fewer problems than the best BUMG method sh ∘ rr ∘ udb.

Looking at the top half of Table 4 (up to difficulty rating of 0.40), the BUMG method based on sh ∘ rr ∘ udb fared best, but for problems more difficult (up to a rating of 0.70) the performance deteriorated and the method crr ∘ udb solved the highest number of problems. For problems with ratings higher than 0.70 SPASS in auto mode solved significantly more problems than the BUMG methods. One problem with rating 1.00 was solved by the crr ∘ udb method (namely, GRP741-1 in 121.86 seconds). Problems in the TPTP library with rating 1.00 have not yet been solved by any other prover.

Table 5 presents an evaluation of the different blocking techniques, listing the number of problems lost and the number of problems gained against the baseline methods in each group. The results confirm the significant positive effect of unrestricted domain blocking for satisfiable problems.

Analysis of the gain and loss of the method based on sh∘rr∘udb against the other methods gave these results: Against rr∘udb 66 problems were gained and

Table 4: Number of problems solved on satisfiable problems, by TPTP problem rating.

Category	#	rr	rr ∘ sdb	rr ∘ udb	rr ∘ spb	rr ∘ upb	sh ∘ rr	sh ∘ rr ∘ sdb	sh ∘ rr ∘ udb	sh ∘ rr ∘ spb	sh ∘ rr ∘ upb	crr	crr ∘ sdb	crr ∘ udb	crr ∘ spb	crr ∘ upb	sh ∘ crr	sh ∘ crr ∘ sdb	sh ∘ crr ∘ udb	sh ∘ crr ∘ spb	sh ∘ crr ∘ upb	auto
0.00	371	112	268	**286**	256	184	154	279	**293**	267	223	80	264	**276**	241	150	80	250	**270**	227	150	231
(0.00, 0.10]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
(0.10, 0.20]	32	1	20	**21**	20	2	6	26	**27**	**27**	7	-	**21**	19	20	1	-	18	**19**	18	-	**29**
(0.20, 0.30]	132	57	**66**	65	59	59	59	**88**	**88**	81	59	54	64	**66**	56	54	55	65	**66**	57	56	39
(0.30, 0.40]	52	4	**6**	**6**	4	4	13	21	**22**	20	13	4	5	**6**	4	4	4	4	**6**	4	4	20
(0.40, 0.50]	125	2	21	**23**	6	7	3	14	**22**	5	8	2	26	**42**	3	6	3	23	**40**	4	6	18
(0.50, 0.60]	27	-	6	**7**	-	-	1	1	**2**	**2**	1	-	3	**7**	-	-	-	-	**1**	-	-	6
(0.60, 0.70]	76	-	1	**2**	-	-	-	1	**2**	-	-	-	5	**8**	-	-	-	**6**	**6**	-	-	-
(0.70, 0.80]	83	-	-	**1**	-	-	-	**1**	**1**	-	-	-	-	-	-	-	-	-	-	-	-	**59**
(0.80, 0.90]	78	-	**1**	-	-	-	-	-	-	-	-	-	-	**4**	-	-	-	1	**2**	-	-	**15**
(0.90, 1.00]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
1.00	1017†	-	-	-	-	-	-	-	-	-	-	-	-	**1**	-	-	-	-	-	-	-	-
Total	-	176	388	**410**	345	256	236	430	**456**	402	311	140	389	**430**	324	215	142	368	**411**	310	216	417

Note: †1017 = 150 SAT + 28 OPN + 839 UNK

Table 5: Evaluation of blocking techniques.

Satisfiable	Baseline	sdb	udb	spb	upb
	rr	-0 / +212	**-4 / +238**	-0 / +169	-1 / +81
	sh ∘ rr	-1 / +195	**-5 / +225**	-2 / +168	-2 / +77
	crr	-0 / +249	**-4 / +294**	-0 / +184	-3 / +78
	sh ∘ crr	-0 / +226	**-5 / +274**	-0 / +168	-3 / +77
Unsatisfiable					
	rr	-211 / +78	-315 / **+83**	-100 / +61	**-77** / +37
	sh ∘ rr	-188 / +106	-225 / **+126**	-81 / +87	**-34** / +52
	crr	-65 / +170	-105 / **+242**	-32 / +78	**-26** / +24
	sh ∘ crr	-52 / +163	-98 / **+190**	-30 / +57	**-16** / +15

20 problems lost; against sh ∘ crr ∘ udb the gain/loss was +90/-45 and against crr ∘ udb it was +85/-59. This non-uniformity suggests each variation of range-restriction had the potential to solve some problems not solvable within the time limit by sh ∘ rr with unrestricted blocking. The biggest variation was against SPASS in auto mode, where 169 problems were gained and 130 problems were lost.

Table 6 displays how many problems were uniquely solved. The first row lists how many problems were uniquely solved over all methods including SPASS in auto mode. Although two of the BUMG methods with unrestricted domain blocking fared better than SPASS in auto mode, the latter solved a significant number of problems that none of the BUMG could solve (namely, 115 problems, or 27.5% of the problems solved by SPASS in auto mode, or 10.2% of all satisfiable problems). This reflected the orthogonality of the underlying methods. Analogously, the relatively low number of problems uniquely solved by the BUMG methods (21 problems, i.e., 1.9% of all satisfiable problems), which is also apparent from the number of problems solved uniquely among the BUMG methods in the second row (25 or 2.2% of all satisfiable problems) can be attributed to the similarity of the underlying methods. An analysis of the number of uniquely solved problems per group of BUMG methods in the third row of the table highlighted the importance of unrestricted domain blocking. While overall no problems were only solved with unary predicate blocking techniques, within the groups there were four problems solved only with unary predicate blocking.

Table 7 gives an impression of the increase in the size of the input files caused by the transformations. Although the file sizes were measured after all comments and white space were removed, variations is name lengths distort the values slightly (which can be seen in the values for shifting). The results therefore need to be interpreted cautiously. The average increase in file size does show a significant effect on the size of the problem for the new range-restriction transformations and also subterm blocking (both subterm domain blocking and subterm predicate blocking). The largest increase in size was observed for the problem SYO600-1 (13.7 fold increase), which contained 380 predicate symbols with arity up to 64, 2 constants and no non-constant function symbols. The main cause for this increase was the large number of clauses added in Step (4)

Table 6: Uniquely solved problems.

Satisfiable	rr	rr ∘ sdb	rr ∘ udb	rr ∘ spb	rr ∘ upb	Total	sh ∘ rr	sh ∘ rr ∘ sdb	sh ∘ rr ∘ udb	sh ∘ rr ∘ spb	sh ∘ rr ∘ upb	Total	crr	crr ∘ sdb	crr ∘ udb	crr ∘ spb	crr ∘ upb	Total	sh ∘ crr	sh ∘ crr ∘ sdb	sh ∘ crr ∘ udb	sh ∘ crr ∘ spb	sh ∘ crr ∘ upb	Total	auto	Total
All methods	-	1	-	-	-		-	-	4	-	-		-	-	11	-	-		-	4	1	-	-		**115**	136
All BUMG methods	-	1	2	-	-		-	-	6	-	-		-	-	**11**	-	-		-	4	1	-	-			25
All BUMG, by group	-	12	**28**	-	-	40	-	4	**25**	1	2	32	-	10	**57**	-	-	67	-	12	**60**	-	1	73		
Unsatisfiable																										
All methods	-	-	1	-	-		-	-	-	-	-		-	-	-	-	-		-	-	-	-	-		**1779**	1780
All BUMG methods	**17**	4	5	2	2		3	5	3	-	3		-	3	3	2	-		-	1	3	-	-			56
All BUMG, by group	**32**	15	18	23	4	92	16	24	**41**	29	14	124	4	23	**90**	10	3	130	8	31	**71**	6	2	118		

Table 7: Average increase in the size of input files.

Method	Avg	Method	Avg	Method	Avg
rr	1.8	crr	1.1	auto	0
rr ∘ sdb	2.4	crr ∘ sdb	1.7		
rr ∘ udb	1.9	crr ∘ udb	1.2		
rr ∘ spb	2.4	crr ∘ spb	1.7		
rr ∘ upb	1.9	crr ∘ upb	1.2		
sh ∘ rr	1.7	sh ∘ crr	1.2		
sh ∘ rr ∘ sdb	2.3	sh ∘ crr ∘ sdb	1.9		
sh ∘ rr ∘ udb	1.8	sh ∘ crr ∘ udb	1.3		
sh ∘ rr ∘ spb	2.3	sh ∘ crr ∘ spb	1.8		
sh ∘ rr ∘ upb	1.7	sh ∘ crr ∘ upb	1.3		

of the rr transformation. For each of the 284 predicate symbols with arity 64 in the problem, 64 clauses were added in Step (4). This is a large number. In contrast for the crr transformations the increase in size was negligible, and also, generally, it was significantly lower. Despite its positive virtues this shows a downside of the rr transformation. For problems containing a large number of function symbols with high arity, Step (5) similarly adds many clauses, even though the transformation overall is still effective.

Analysis of the problems solved without any form of blocking revealed a large number belonged to the Bernays-Schönfinkel class: 131/176 (74%) for rr, 132/236 (56%) for sh ∘ rr, 133/140 (95%) for crr, and 134/142 (94%) for sh ∘ rr. These results confirmed the expectation that more problems not solvable with the crr transformation can be solved with the rr transformation and the benefits of reducing the number of terms created.

Although the main purpose of BUMG methods is *disproving* theorems and generating models for satisfiable problems, for completeness we report in Tables 8 and 9 the results for unsatisfiable clausal TPTP problems. The results were not as uniform as for satisfiable problems. However some general observations can be made. SPASS in auto mode fared best overall, and did so in all TPTP categories and each problem rating category. For unsatisfiable problems the drawback of BUMG methods is that clauses need to be exhaustively grounded and each branch in the derivation tree needs to be closed. The dominance of SPASS in auto mode is thus not surprising.

For the BUMG methods, a general deterioration in performance could be observed for shifting, when comparing the results for the groups with baselines sh ∘ rr and sh ∘ crr to the respective groups without shifting. This is plausible because shifting leads to fewer negative literals in clauses and more positive literals thus reducing the constraining effect and leading to more splitting. For problems with higher rating, shifting did seem to have a positive effect; for instance, in the $(0.40, 0.50]$ range, sh ∘ rr solved 70 problems whereas rr solved 32 problems.

Within the BUMG groups we expected best performance for the baseline transformations, because these do not involve blocking and performing many blocking steps lead to a significant overhead. However only for the first group the rr transformation fared best. In combination with classical range-restriction crr, somewhat surprisingly, the best results were obtained with unrestricted do-

main blocking, the most expensive form of blocking, because it is applicable to any terms. Among the blocking techniques in each case the highest gain was obtained for unrestricted domain blocking (see Table 5). However also the greatest loss was observed for this blocking technique. The smallest loss and lowest gain was obtained for upb blocking. The high loss for udb could be a reflection of the high increase in splitting steps preventing quicker detection of contradictions. Analogously the small loss for upb could be attributable to the smallest number of additional splitting steps among the blocking techniques. The high gain for udb blocking suggests the inference process panned out significantly differently leading to solutions not found with the other techniques. This seems to be supported by the results in the third row of Table 6 according to which, with one exception, the largest number of uniquely solved problems in each group was obtained with udb blocking. The exception was the first group, where rr led to the largest number of uniquely solved problems. Among all the BUMG methods, rr solved the largest number of problems not solved by any of the other methods. However, these results pale against the number of uniquely solved problems by SPASS in auto mode. Only one problem was solved by a BUMG method which was not solved by SPASS in auto mode.

9.2 Findings

Several findings can be drawn from the results. The results have confirmed our expectation that unrestricted domain blocking is a powerful technique, which helps discover finite models more often than with the other blocking techniques. The results suggest the technique is indispensable for bottom-up model generation. Both in combination with the new range-restricting transformation, and the classical range-restricting transformation, good results have been obtained. Overall, the method based on new range-restriction, shifting and unrestricted domain blocking performed best on the sample. On satisfiable problems with higher difficulty rating this method was however gradually edged out by the method based on classical range-restriction and unrestricted domain blocking. This suggests there is a trade-off between the rr transformation, which is based on a non-trivial transformation but does restrict the creation of terms, and the simpler crr transformation, which has to rely on blocking to restrict the creation of terms.

The results for subterm domain blocking were good and often not far behind unrestricted domain blocking for satisfiable problems. In contrast, predicate blocking seems not to be effective on many problems. We attribute this to the nature of the problems in the TPTP library.

An investigation with SPASS-Yarralumla on translations of modal logic problems has revealed a different picture [Schmidt et al., 2014]. There, the best performance was obtained with subterm domain blocking for both satisfiable and unsatisfiable problems. Better results than for unrestricted domain blocking were also obtained with subterm predicate blocking and unrestricted predicate blocking. Better performances for subterm and predicate blocking are also expected on problems stemming from (cyclic) description logic knowledge bases. Experiments with blocking restricted by excluding a finite subset of the domain have shown better results than for unrestricted domain blocking for consistency testing on a large corpus of ontologies [Khodadadi et al., 2013]. The better performance for restricted forms of blocking on modal and description

Table 8: Number of problems solved on unsatisfiable clausal TPTP problems.

Category	#	rr	rr ∘ sdb	rr ∘ udb	rr ∘ spb	rr ∘ upb	sh ∘ rr	sh ∘ rr ∘ sdb	sh ∘ rr ∘ udb	sh ∘ rr ∘ spb	sh ∘ rr ∘ upb	crr	crr ∘ sdb	crr ∘ udb	crr ∘ spb	crr ∘ upb	sh ∘ crr	sh ∘ crr ∘ sdb	sh ∘ crr ∘ udb	sh ∘ crr ∘ spb	sh ∘ crr ∘ upb	auto
ALG	133	**13**	10	11	10	12	**7**	6	**7**	6	6	18	18	**20**	19	19	5	5	4	**6**	4	**62**
ANA	83	7	**11**	**11**	10	7	6	**7**	**7**	6	6	4	7	**10**	6	4	3	4	**5**	3	3	**40**
BOO	74	38	32	28	**43**	38	38	31	26	**42**	37	33	36	**38**	33	33	33	**36**	34	33	33	**62**
CAT	52	**38**	33	33	34	**38**	29	31	**32**	29	28	34	**36**	35	35	34	27	**30**	**30**	**30**	28	**51**
COL	225	**32**	23	22	25	28	19	18	16	17	**19**	**32**	31	25	**32**	31	8	**15**	**15**	8	8	**149**
COM	14	6	6	6	6	6	**6**	**6**	5	**6**	**6**	5	5	5	5	5	4	4	4	4	4	**12**
FLD	175	24	39	44	38	**48**	19	38	**45**	38	43	23	32	**36**	35	19	17	29	**37**	29	17	**99**
GEO	187	**68**	**68**	67	67	66	71	**74**	66	72	71	45	43	**53**	45	42	45	46	**47**	44	44	**121**
GRA	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	**1**
GRP	798	233	183	171	**240**	225	**307**	231	260	291	298	319	378	**424**	337	322	358	385	**406**	353	357	**673**
HEN	64	**41**	**41**	34	40	**41**	34	**37**	32	34	34	31	**42**	40	31	29	36	**43**	33	35	36	**62**
HWC	4	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	1	1	1	1	1	**2**
HWV	116	**44**	39	34	40	35	21	21	21	**23**	20	18	18	**20**	18	19	**15**	14	10	13	12	**72**
KRS	17	9	9	9	9	9	9	9	9	9	9	7	**9**	**9**	**9**	**9**	5	8	**9**	8	8	**9**
LAT	250	39	33	29	**40**	39	39	25	24	**43**	37	56	56	**64**	57	56	70	65	56	**71**	70	**91**
LCL	625	**44**	40	35	39	45	**19**	**19**	17	**19**	**19**	43	**44**	**44**	**44**	43	21	24	**26**	23	21	**336**
LDA	23	**3**	2	2	**3**	**3**	**3**	**3**	2	**3**	**3**	5	5	5	5	5	5	**6**	5	5	5	**6**
MGT	67	**41**	38	35	37	42	**52**	40	29	40	46	18	**30**	25	21	20	12	**23**	17	13	12	**61**
MSC	20	9	**11**	9	10	9	8	**10**	7	9	7	8	**10**	**10**	9	8	7	**10**	8	7	7	**14**
NLP	22	17	**18**	**18**	**18**	17	**20**	19	18	18	**20**	9	**18**	14	**18**	9	9	**18**	14	**18**	9	**22**
NUM	86	**15**	10	10	13	13	10	8	9	9	**11**	10	9	8	10	**11**	**8**	7	**8**	**8**	**8**	**36**
PLA	51	4	**5**	4	4	**5**	3	**4**	3	3	**4**	3	**4**	2	**4**	3	2	**3**	1	**3**	**3**	**31**
PUZ	72	**46**	45	37	45	45	**38**	37	36	37	**38**	**44**	43	35	43	41	**34**	**34**	30	33	33	**55**
REL	107	**4**	3	2	**4**	**4**	3	2	2	**4**	3	5	**7**	6	5	5	5	**6**	**6**	5	5	**37**
RNG	83	20	11	8	**23**	19	19	10	9	**22**	19	10	**14**	**14**	11	10	10	**14**	**14**	11	11	**48**
ROB	31	**4**	3	2	**4**	**4**	**5**	3	1	**5**	**5**	3	**5**	3	3	3	3	**6**	4	5	3	**14**
SCT	98	8	8	8	8	8	2	2	2	2	2	11	11	11	11	11	2	2	2	2	2	**27**
SET	450	**68**	63	57	58	66	46	50	**52**	47	45	53	50	51	**54**	53	29	36	**40**	29	28	**290**
SWC	383	**120**	107	76	100	81	75	74	75	74	**89**	**111**	106	100	105	110	99	**100**	89	99	97	**281**
SWV	855	100	72	70	92	**104**	43	45	38	**48**	45	**77**	72	71	**77**	**77**	29	31	**33**	29	28	**373**
SWW	33	-	-	-	-	-	-	-	-	-	-	1	1	1	1	1	-	-	-	-	-	**12**
SYN	621	**330**	329	321	326	328	318	**326**	318	319	317	321	321	314	319	**322**	300	**307**	**306**	300	303	**508**
SYO	17	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	-	1	-	1	1	**3**
TOP	5	3	3	3	3	3	1	2	**3**	1	1	2	**3**	**3**	**3**	**3**	1	1	1	1	1	**5**
Total	5842	**1431**	1298	1199	1392	1391	1273	1191	1174	1279	**1291**	1363	1468	**1500**	1409	1361	1204	**1315**	1296	1231	1203	**3665**

Table 9: Result summary wrt. problem rating on unsatisfiable clausal TPTP problems.

Category	#	rr	rr ∘ sdb	rr ∘ udb	rr ∘ spb	rr ∘ upb	sh ∘ rr	sh ∘ rr ∘ sdb	sh ∘ rr ∘ udb	sh ∘ rr ∘ spb	sh ∘ rr ∘ upb	crr	crr ∘ sdb	crr ∘ udb	crr ∘ spb	crr ∘ upb	sh ∘ crr	sh ∘ crr ∘ sdb	sh ∘ crr ∘ udb	sh ∘ crr ∘ spb	sh ∘ crr ∘ upb	auto
0.00	1409	871	830	796	860	**885**	734	745	722	744	**746**	739	**769**	747	756	737	590	**654**	647	614	597	**1360**
(0.00, 0.10]	480	**244**	203	187	236	239	210	184	167	**216**	210	236	239	**245**	237	234	204	**213**	212	202	201	**477**
(0.10, 0.20]	594	**151**	145	118	148	137	133	123	112	140	**141**	151	**165**	162	164	150	145	**157**	144	**157**	144	**529**
(0.20, 0.30]	507	**98**	85	72	87	79	70	61	63	65	**72**	85	89	**94**	85	83	67	70	**71**	66	64	**382**
(0.30, 0.40]	258	**30**	24	18	**30**	24	38	30	**40**	35	36	42	55	**61**	51	44	51	**53**	46	49	51	**226**
(0.40, 0.50]	513	**32**	11	6	24	22	**70**	45	53	62	69	75	116	**139**	86	76	89	112	**125**	89	88	**344**
(0.50, 0.60]	207	3	-	1	**4**	2	13	2	**14**	9	13	22	23	**30**	20	23	27	**30**	**30**	26	27	**122**
(0.60, 0.70]	298	2	-	1	3	3	**4**	1	3	**6**	4	12	11	**13**	9	12	**19**	14	11	16	**19**	**110**
(0.70, 0.80]	324	-	-	-	-	-	1	-	-	**2**	-	1	1	**6**	1	2	**4**	**4**	**4**	**4**	**4**	**85**
(0.80, 0.90]	344	-	-	-	-	-	-	-	-	-	-	-	-	**1**	-	-	**1**	**1**	**1**	**1**	**1**	**25**
(0.90, 1.00]	259	-	-	-	-	-	-	-	-	-	-	-	**2**	-	-	-	**7**	**7**	5	**7**	**7**	**5**
1.00	1516	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Total	-	**1431**	1298	1199	1392	1391	1273	1191	1174	1279	**1291**	1363	1468	**1500**	1409	1361	1204	**1315**	1296	1231	1203	**3665**

Note: †1516 = 649 UNS + 28 OPN + 839 UNK

logic problems can be attributed to mainstream modal and description logics having the finite tree model property. This means every satisfiable formula holds in a model based on a finite tree, which is not a property of first-order formulae.

The results showed BUMG methods were good for disproving theorems and generating models for satisfiable problems. For unsatisfiable problems BUMG methods were however significantly less efficient than SPASS in auto mode. For theorem proving purposes a limitation of BUMG methods is that they require full grounding. It can be seen already from very small unsatisfiable examples that a complete BUMG derivation tree can be very large, whereas resolution proofs are significantly shorter.

Compared to resolution, an advantage of BUMG methods for satisfiable problems is the division of the search space into branches which are individually constructed and individually processed. As a consequence, if the right decisions are made at branching points models can be found more quickly. When the branching point decisions are less optimal the performance can deteriorate dramatically, particularly if the search is trapped in a branch with only infinite models. This could be another explanation for the lower success rate of the BUMG methods observed for more difficult satisfiable problems. For problems where only infinite models exist, clearly other methods are better.

10 Conclusions

We have presented and tested a number of enhancements for BUMG methods. An important aspect is that our enhancements exploit the strengths of readily available BUMG system with only modest modifications. Our range restriction technique is a refinement of existing transformations to range-restricted clauses in that terms are added to the domain of interpretation on a 'by need' basis. Moreover, we have presented methods that allow us to extend BUMG methods with blocking techniques related to loop checking techniques with a long history in the more specialized setting of modal and description logics.

The experimental evaluation has shown blocking techniques are indispensable in BUMG methods for satisfiable problems. In particular, unrestricted domain blocking turned out to be the most powerful technique on problems from the TPTP library. Limiting the creation of terms during the inference process by using the new range restricting transformation paid off, leading to better results. It is particularly advisable together with the shifting transformation. The experimental results however also show that classical range restriction together with unrestricted blocking is a good complementary method. Because model generation methods are not just aimed at showing the existence of models but are built to construct and return models, when no models exists the entire search space must be traversed, which has led to inferior performance compared to saturation-based resolution.

Our bottom-up model generation approach is especially suitable for generating small models and it is possible to show the approach using unrestricted domain blocking allows us to compute finite models when they exist. The models produced by subterm blocking and predicate blocking are not as small as those produced by unrestricted domain blocking. In particular, the generated models do not need to be Herbrand models. It follows from how the transformations work that the generated models are quasi-Herbrand models, in the following

sense. Whenever $\mathsf{dom}(s)$ and $\mathsf{dom}(t)$ hold in the (Herbrand) model constructed by the BUMG method, then (as in Herbrand interpretations) the terms s and t are mapped to themselves in the associated (possibly non-Herbrand) model. Reconsidering the example in the Introduction of the two unit clauses $\mathsf{P}(\mathsf{a})$ and $\mathsf{Q}(\mathsf{b})$, the associated model maps a and b to themselves, regardless as to which transformations are applied as long as it includes a form of subterm blocking. In this way, more informative models are produced than those computed by, for example, MACE- and SEM-style finite model searchers (and also unrestricted domain blocking). From an applications perspective, this can be an advantage because larger models are more likely to be helpful to a user debugging mistakes in the formal specification of a program or protocol, or an ontology engineer trying to discover why an expected entailment does not follow from an ontology.

Research in automated theorem proving on developing decision procedures has concentrated on developing refinements of resolution, mainly ordering refinements, for deciding solvable fragments of first-order logic. Fragments decidable with ordered resolution are complementary to the fragments that can be decided by refinements using the techniques presented in this paper. We have thus extended the set of techniques available for resolution methods to turn them into more effective and efficient (terminating) automated reasoning methods. In particular, we have shown that all procedures based on hyperresolution, or BUMG methods, can decide the Bernays-Schönfinkel class and the class of BS clauses with equality.

Studying how well the ideas and techniques discussed in this paper can be exploited and behave in dedicated BUMG provers, tableau-based provers and other provers (including resolution-based provers) is very important but is beyond the scope of the present paper. Initial results with another prover, Darwin [Baumgartner et al., 2006], are very encouraging. An in-depth comparison and analysis of BUMG approaches with our techniques and MACE-style or SEM-style model generation would also be of interest. Another source for future work is to combine the presented transformations with other BUMG techniques, such as magic sets transformations [Hasegawa et al., 1997, Stickel, 1994], a typed version of range-restriction [Baumgartner et al., 1997], and minimal model computation [Bry and Yahya, 2000, Bry and Torge, 1998, Papacchini and Schmidt, 2011]. Having been designed to be generic, we believe that our transformations carry over to formalisms with default negation, which could provide a possible basis for enhancements to answer-set programming systems.

Acknowledgements. The second author is grateful to Christoph Weidenbach and Uwe Waldmann for hosting her during 2010 and 2013–2014. In this time the implementation of SPASS-yarralumla was completed and the experimental evaluation was undertaken on the cluster of the Max-Planck-Institut für Informatik, Saarbrücken. We thank Uli Furbach, Dmitry Tishkovsky, Uwe Waldmann and Christoph Weidenbach for useful discussions and comments on this research. This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) (grants EP/F068530/1 and EP/H043748/1), NICTA, Canberra, Australia, and the Max-Planck-Institut, Saarbrücken, Germany.

References

F. Baader and U. Sattler. An overview of tableau algorithms for description logics. *Studia Logica*, 69:5–40, 2001.

L. Bachmair and H. Ganzinger. Equational reasoning in saturation-based theorem proving. In W. Bibel and P. H. Schmitt, editors, *Automated Deduction—A Basis for Applications*, pages 353–397. Kluwer, 1998.

L. Bachmair and H. Ganzinger. Resolution theorem proving. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, pages 19–99. Elsevier, 2001.

P. Baumgartner and R. A. Schmidt. Blocking and other enhancements for bottom-up model generation methods. In U. Furbach and N. Shankar, editors, *Automated Reasoning: IJCAR 2006*, volume 4130 of *Lecture Notes in Artificial Intelligence*, pages 125–139. Springer, 2006.

P. Baumgartner, U. Furbach, and I. Niemelä. Hyper tableaux. In J. J. Alferes, L. M. Pereira, and E. Orlowska, editors, *Logics in Artificial Intelligence: JELIA'96*, volume 1126 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 1996.

P. Baumgartner, U. Furbach, and F. Stolzenburg. Computing answers with model elimination. *Artificial Intelligence*, 90(1–2):135–176, 1997.

P. Baumgartner, A. Fuchs, and C. Tinelli. Implementing the model evolution calculus. *International Journal of Artificial Intelligence Tools*, 15(1):21–52, 2006.

P. Baumgartner, U. Furbach, and B. Pelzer. Hyper tableaux with equality. In F. Pfenning, editor, *Automated Deduction: CADE-21*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 492–507. Springer, 2007.

M. Bezem. Disproving distributivity in lattices using geometry logic. In *Proceedings of CADE-20 Workshop on Disproving*, 2005.

M. P. Bonacina, C. Lynch, and L. M. de Moura. On deciding satisfiability by theorem proving with speculative inferences. *Journal of Automated Reasoning*, 47(2):161–189, 2011.

F. Bry and S. Torge. A deduction method complete for refutation and finite satisfiability. In J. Dix, L. Fariñas del Cerro, and U. Furbach, editors, *Logics in Artificial Intelligence: JELIA'98*, volume 1489 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 1998.

F. Bry and A. Yahya. Positive unit hyperresolution tableaux for minimal model generation. *Journal of Automated Reasoning*, 25(1):35–82, 2000.

K. Claessen and N. Sörensson. New techniques that improve MACE-style finite model building. In P. Baumgartner and C. G. Fermüller, editors, *Proceedings of CADE-19 Workshop on Model Computation*, 2003.

M. Fujita, J. Slaney, and F. Bennett. Automatic generation of some results in finite algebra. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 52–57. Morgan Kaufmann, 1995.

T. Geisler, S. Panne, and H. Schütz. Satchmo: The compiling and functional variants. *Journal of Automated Reasoning*, 18(2):227–236, 1997.

R. Hasegawa, K. Inoue, Y. Ohta, and M. Koshimura. Non-horn magic sets to incorporate top-down inference into bottom-up theorem proving. In *Automated Deduction: CADE-14*, volume 1249 of *Lecture Notes in Computer Science*, pages 176–190. Springer, 1997.

U. Hustadt and R. A. Schmidt. On the relation of resolution and tableaux proof systems for description logics. In T. Dean, editor, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence: IJCAI'99*, pages 110–115. Morgan Kaufmann, August 1999.

U. Hustadt and R. A. Schmidt. MSPASS: Modal reasoning by translation and first-order resolution. In R. Dyckhoff, editor, *Automated Reasoning with Analytic Tableaux and Related Methods, International Conference: TABLEAUX 2000*, volume 1847 of *Lecture Notes in Artificial Intelligence*, pages 67–71. Springer, 2000.

U. Hustadt, R. A. Schmidt, and C. Weidenbach. MSPASS: Subsumption testing with SPASS. In P. Lambrix, A. Borgida, M. Lenzerini, R. Möller, and P. Patel-Schneider, editors, *Proceedings of International Workshop on Description Logics'99*, pages 136–137. Linköping University, 1999.

M. Khodadadi, R. A. Schmidt, and D. Tishkovsky. A refined tableau calculus with controlled blocking for the description logic $\mathcal{SHOI}$. In D. Galmiche and D. Larchey-Wendling, editors, *Automated Reasoning with Analytic Tableaux and Related Methods: TABLEAUX 2013*, volume 8123 of *Lecture Notes in Computer Science*, pages 188–202. Springer, 2013.

C. Lynch. Unsound theorem proving. In *Computer Science Logic, 18th International Workshop, CSL*, volume 3210 of *LNCS*, pages 473–487. Springer, 2004.

R. Manthey and F. Bry. SATCHMO: a theorem prover implemented in Prolog. In E. Lusk and R. Overbeek, editors, *Automated Deduction: CADE-9*, volume 310 of *Lecture Notes in Computer Science*, pages 415–434. Springer, 1988.

W. McCune. A Davis-Putnam Program and its Application to Finite First-Order Model Search: Quasigroup Existence Problems. Technical Report MCS-TM-194, ANL, 1994.

W. McCune. Mace4 reference manual and guide. Technical Memorandum 264, Argonne National Laboratory, 2003.

R. Nieuwenhuis and A. Rubio. Paramodulation-based theorem proving. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, pages 371–443. Elsevier, 2001.

F. Papacchini and R. A. Schmidt. A tableau calculus for minimal modal model generation. *Electronic Notes in Theoretical Computer Science*, 278(3):159–172, 2011.

B. Pelzer and C. Wernhard. System description: E-KRHyper. In F. Pfenning, editor, *Automated Deduction: CADE-21*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 508–513. Springer, 2007.

A. Riazanov and A. Voronkov. The design and implementation of VAMPIRE. *AI Communications*, 15(2-3):91–110, 2002.

J. A. Robinson. Automatic deduction with hyper-resolution. *International Journal of Computer Mathematics*, 1(3):227–234, 1965.

R. A. Schmidt and U. Hustadt. Solvability with resolution of problems in the Bernays-Schönfinkel class. Presented at Dagstuhl Seminar 05431, 2006, and ARW 2006 in Bristol, 2005.

R. A. Schmidt and D. Tishkovsky. Using tableau to decide expressive description logics with role negation. In K. Aberer, K.-S. Choi, N. Fridman Noy, D. Allemang, K.-I. Lee, L. J. B. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber, and P. Cudré-Mauroux, editors, *The Semantic Web: ISWC 2007 + ASWC 2007*, volume 4825 of *Lecture Notes in Computer Science*, pages 438–451. Springer, 2007.

R. A. Schmidt and D. Tishkovsky. A general tableau method for deciding description logics, modal logics and related first-order fragments. In A. Armando, P. Baumgartner, and G. Dowek, editors, *Automated Reasoning: IJCAR 2008*, volume 5195 of *Lecture Notes in Computer Science*, pages 194–209. Springer, 2008.

R. A. Schmidt and D. Tishkovsky. Automated synthesis of tableau calculi. *Logical Methods in Computer Science*, 7(2):1–32, 2011.

R. A. Schmidt and D. Tishkovsky. Using tableau to decide description logics with full role negation and identity. *ACM Transactions on Computational Logic*, 15(1), 2013.

R. A. Schmidt, J. G. Stell, and D. Rydeheard. Axiomatic and tableau-based reasoning for Kt(H,R). In R. Goré and A. Kurucz, editors, *Advances in Modal Logic, Volume 10*, pages ??–??, London, 2014. College Publications.

S. Schulz. System Description: E 1.8. In K. McMillan, A. Middeldorp, and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning: LPAR 2013*, volume 8312 of *LNCS*, pages 735–743. Springer, 2013.

J. Slaney. FINDER (finite domain enumerator): Notes and guide. Technical Report TR-ARP-1/92, Australian National University, 1992.

R. M. Smullyan. *First Order Logic*. Springer, Berlin, 1971.

M. E. Stickel. Upside-down meta-interpretation of the model elimination theorem-proving procedure for deduction and abduction. *Journal of Automated Reasoning*, 13(2):189–210, 1994.

C. Weidenbach, R. A. Schmidt, T. Hillenbrand, R. Rusev, and D. Topic. System description: SPASS version 3.0. In F. Pfenning, editor, *Automated Deduction: CADE-21*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 514–520. Springer, 2007.

C. Weidenbach, D. Dimova, A. Fietzke, R. Kumar, M. Suda, and P. Wischnewski. SPASS version 3.5. In R. A. Schmidt, editor, *Automated Deduction: CADE-22*, volume 5663 of *Lecture Notes in Computer Science*, pages 140–145. Springer, 2009.

C. Wernhard. System description: KRHyper. In *Proceedings of CADE-19 Workshop on Model Computation*, 2003.

H. Zhang. SEM: A system for enumerating models. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 298–303. Morgan Kaufmann, 1995.